Statistical Model-Based Corneal Reconstruction

by

Justin A. Eichel

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2013

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Precise measurements of corneal layer thickness are required to treat, evaluate risk of, and determine the progression of pathologies within the eye. The thickness measurements are typically acquired as 2D images, known as tomograms, from an optical coherence tomography (OCT) system. With the creation of ultra-high resolution OCT (UHROCT), there is active research in precisely measuring, *in vivo*, previously unresolvable corneal structures at arbitrary locations within the cornea to determine their relationship with corneal health.

In order to obtain arbitrary corneal thickness measurements, existing reconstruction techniques require the cornea to be densely sampled so that a 3D representation can be interpolated from a stack of tomograms. Unfortunately, tomogram alignment relies solely on image properties such as pixel intensity, and does not constrain the reconstruction to corneal anatomy. Further, the reconstruction method cannot properly compensate for eye-motion. The deficiencies due to eye-motion are exacerbated due to the amount of time required in a single imaging session to acquire a sufficient number of tomograms in the region of interest.

The proposed methodology is the first to incorporate models of the anatomy and the imaging system to address the limitations of existing corneal reconstruction methods. By constructing the model in such a way as to decouple anatomy from the imaging system, it becomes less computationally expensive to estimate model parameters. The decoupling provides an iterative methodology that can allow additional constraints to be introduced in the future. By combining sparsely sampled UHROCT measurements with a properly designed corneal model, reconstruction allows researchers to determine corneal layer thicknesses at arbitrary positions in both sampled and unsampled regions.

The proposed methodology demonstrates an approach to decouple anatomy and physiology from measurements of a cornea, allowing for characterization of pathologies through corneal thickness measurements. Another significant contribution resulting from the corneal model allows five of the corneal layer boundaries to be automatically located and has already been used to process thousands of UHROCT tomograms. Recent studies using this method have also been used to correlate contact-lens wear to hypoxia and corneal layer swelling. While corneal reconstruction represents the main application of this work, the reconstruction methodology can be extended to other medical imaging domains and can even represent temporal changes in tissue with minor modifications to the framework.

Acknowledgements

Contents

List of Figures

List of Tables

Abbreviations

Nomenclature

This section contains reference to all major notation, categorized by primary mathematical purpose. The following general rules are applied throughout the document.

1. Capital letters are used to represent matrices and sets.

2. Capital script letters are used to denote kernels.

3. Vectors are denoted in bold, for example $\mathbf{v}$.

4. Directional vectors often used to represent basis vectors are represented with an overline, for example $\overrightarrow{\mathbf{d}}$.

5. Elements in a vector, matrix, or set, are indexed using subscripts, for example v_0, $M_{0,3}$, and S_i respectively.

6. The symbol : in a subscript indicates all elements in the dimension where the symbol is used, for example $M_{:,1}$ is a column vector containing all elements in the first column of matrix M and $M_{3,:}$ is a row vector containing all the elements in row three of matrix M.

7. Scalars are never bold.

8. Functions are represented by either a capital or lower case letter with optional subscripts followed by parentheses containing parameters, $f(a, b, c)$.

9. To avoid confusion with the function notation, multiplication involving parentheses is explicit, for example $a \cdot (b + c)$.

10. 2D coordinates are represented in the form (x, y) and 3D coordinates are represented in the form (x, y, z).

Statistical reconstruction

Measurements

Random fields

Markov random fields

Deterministic random field

Decoupling

Forward model

Reconstruction

Tomogram pose estimation

Corneal imaging

Tomograms

Imaging planes

Coordinate transformation

Corneal boundary localization

Corneal reconstruction

Structure

Motion

Reconstruction

Chapter 1

Introduction

Ralph Waldo Emerson wrote that "science does not know its debt to imagination," [37]. While problem solving requires a great deal of technical knowledge and skill, it also requires the creativity and persistence to shape that knowledge, transforming it to make the world a better place. Statistical model-based reconstruction encompasses these ideals. The surrounding world can be expressed as a mathematical model through creativity, knowledge of structure, and knowledge of behaviour. Using the model as context, observations are given meaning, which leads to better understanding and more accurate models. Then, it is up to humankind to use that knowledge to improve the world.

1.1 Statistical model-based reconstruction

When statistical model-based reconstruction is applied to a discipline, it helps satisfy a number of goals. It can provide better insight from observations, provide a means to judge observation accuracy, and a means to estimate unknown parameters even if direct measurements are unavailable. The advantages of applying statistical model-based reconstruction to the application of corneal reconstruction are presented here.

Often researchers are not interested in raw measurements, but are interested in the implications of the raw measurements. During reconstruction, these raw measurements are fitted to a statistical model. If the model was purposely designed, the model parameters, estimated from the raw measurements, can provide better insight into the measurement implications. For example, the length of the cranium has much more meaning when it is combined with an anthropological model allowing the volume to imply sex [1].

A statistical model also provides a way to determine measurement accuracy. When the model parameters are estimated during reconstruction, measurement obtained from the real world will differ from measurements sampled from the reconstruction due to measurement uncertainty. The difference between real-world and reconstructed measurements are referred to as residuals. Statistical anomalies can be detected from these residuals indicating the inability of the model to capture a trend, often in the form of unpredictable measurement uncertainty.

Statistical model-based reconstruction can generate estimates of regions that have not been measured directly. After model parameters are estimated using the available measurements, estimation methods can be applied to unsampled regions. Statistical modelling ensures that relationships between measurements are mathematically formulated so that unsampled measurements can be derived from a model. The mathematical model should incorporate application specific knowledge, including available measurements of the subject. For example, current raw wind and pressure measurements can be combined with a formally developed weather forecasting model to predict future conditions [64] or, for a spatial example, unsampled regions of a cornea can be estimated from a few sparse measurements.

1.2 Corneal reconstruction

Recent advances in corneal imaging allow researchers to visualize corneal structures as thin as 5 μm at a rate of 29 to 47 tomograms per second from patients without paralytics and without coming in contact with the eye [18, 86]. A tomogram is high resolution images of the cornea. Although imaging has advanced, existing corneal reconstruction techniques still require every voxel, a point in a 3D volume, to be directly sampled [79, 106, 114, 134, 148] in order for the researcher to obtain and use a 3D visualization of the data. Sampling every voxel requires millions of measurements. Further, eye-motion during imaging makes it difficult to determine exactly where the imaging system sampled the eye [140].

A statistical corneal model-based reconstruction can allow researchers to quantify and reduce measurement uncertainty and to estimate corneal structures in unsampled regions based on known corneal anatomy and physiology. Using historic data statistical corneal model-based reconstruction can be improved and extended to provide a comprehensive corneal model. Historic patient data can be registered to the model allowing treatment and diagnostic measurements to be consistently measured at arbitrary places within the cornea.

While reserving a detailed discussion of corneal anatomy, imaging systems, and reconstruction limitations for Section 2.2, a high level overview is presented here so that existing limitations and challenges of corneal reconstruction can be better understood.

Ophthalmological overview

The cornea is the outermost structure in the eye, protecting the interior and refracting light into the lens where the light is ultimately focused on the retina [125, 141]. There are six distinct layers within the cornea that are responsible for its structure and transparency: epithelium, Bowman's layer, stroma, Dua's layer, Decement's membrane, and endothelium [34, 50, 81, 111]. Pathologies cause the six layers to change shape, which can decrease the cornea's transparency and cause light entering the cornea to no longer project at the proper location on the retina [106, 133]. Diseases also result in dynamic alterations in the thickness of some or all of the corneal layers [73, 145]. Therefore, precise measurements of the thicknesses of the corneal layers, as a function of spatial location in the cornea and over time, can be used as a biomarker for the clinical investigation at the onset and during stages of development of various corneal diseases [25, 26, 51, 62, 73, 82].

Optical coherence tomography

Optical coherence tomography (OCT) is an imaging method for non-invasive, high resolution ($1-15\ \mu m$), volumetric imaging of biological tissue at depths of one to two millimeters below the tissue surface [41, 56]. Unlike other imaging systems, OCT does not require contact with the cornea, reducing the risk of eye damage and infection. Within the last decade, Fourier-domain optical coherence tomography (FDOCT) has been established as an important clinical diagnostic modality in ophthalmology for non-invasive clinical assessment of retinal and corneal health [32, 103, 122]. FDOCT has the advantage of improved scanning times since measurements in the depth direction are sampled simultaneously using Fourier transforms instead of being individually sampled using a moving apparatus [40].

Corneal thickness measurements

Until recently, the axial resolution of OCT technology in the human cornea was limited to approximately 10 μm to 15 μm, which is insufficient for resolving the thin corneal layers, such as Bowman's membrane or the Descemet's endothelium complex. Consequently, segmentation algorithms were only capable of measuring the thicker epithelium and stroma layers. For example, Benchmann [8] measures the thickness of these layers, but that

algorithm can only measure near the apex of the cornea. Fishman developed a method that requires the user to manually select points on the image allowing measurements between these points, a time consuming and potentially inconsistent process [43]. Wang et al.'s algorithm measured the layer thickness for the epithelial layer [138]. Due to the large refractive index difference, the boundaries of these layers appear to have high contrast in the OCT images, therefore Wang et al.'s segmentation algorithm was able to evaluate the thickness of this layer. However, Wang et al.'s [138] method fails to locate the inner corneal layers, Bowman's layer, stroma, and Descemet's membrane, since they have a lower contrast. An alternative localization method is required to segment these layers.

There is also considerable interest in constructing 3D corneal representations from 2D segmented images, in which case hundreds of OCT images can be segmented in order to measure thickness at arbitrary positions within a 3D corneal representation. Without model-based reconstruction to estimate corneal structure from unsegmented regions, a $1000 \times 1000 \times 512$ 3D model would require 1000 tomograms to be acquired and segmented. A model-based reconstruction can infer structure from OCT measurements requiring fewer tomograms and allowing structure estimation in unsampled corneal regions. An automated algorithm to locate each corneal layer is essential to build a structural model from corneal layer statistics collected from thousands of tomograms, since a trained technician can require approximately 15 to 20 minutes to manually segment a single 2D UHROCT image [35]. Further, the boundaries located using an automated boundary localization algorithm can be used to estimate corneal structure during reconstruction.

1.3 Thesis scope

The previous sections have motivated the need to apply statistical model-based reconstruction to corneal imaging. A model-based corneal reconstruction can allow researchers to take measurements from arbitrary locations without having to sample each point within the reconstruction. Further, a corneal model can allow OCT error to be quantified using residuals from the reconstruction process. By applying anatomical corneal model constraints to the measurement data, measurements taken from the reconstruction should be more consistent with corneal anatomy.

To achieve the goal of a corneal reconstruction, detailed background is provided in Chapter 2, a problem formulation is presented in Chapter 3, and the following high level objectives are required to create the first model-based corneal reconstruction method that integrates both anatomy and raw imaging data.

1. **Extend general reconstruction theory**

   To allow statistical model-based reconstruction to be applied to corneal imaging and other similar applications, the general reconstruction theory detailed in Section 2.1 is expanded in Chapter 4. This extension illustrates how measurements collected within imaging planes can be mapped into a random field and presents a proposed method allowing a corneal model to be decoupled into anatomy and scattering potential data. This extension is necessary since this is the first time that this reconstruction theory is applied to OCT corneal data.

2. **Create 2D corneal structure model**

   As a first step toward understanding and designing a corneal structure model, a statistical analysis of corneal layer structure is performed through the creation of an automated corneal layer boundary algorithm in Chapter 5. Then structures from 2D tomogram data can be analyzed. This model based method is the first to automatically segment five corneal layers from OCT.

3. **Create 3D corneal model**

   Using the 2D statistical model created in Chapter 5 a 3D structural model can be designed in Section 6.1 and a model relating OCT measurements can be designed in Section 6.2. Then using background from Subsection 2.2.3, a corneal eye-motion model can be designed in Section 6.3 to help reduce measurement location error. This objective is fundamental to corneal reconstruction since the corneal model is used to estimate unknown measurements and correct noisy data. The anatomical portion of the model is the first to incorporate corneal layer boundaries as part of the reconstruction process and can be extended in the future to include additional structures as desired.

4. **Generate synthetic cornea from model**

   Since 3D corneal OCT samples from healthy human subjects that have been manually segmented and manually registered in 3D space are not available in datasets, and synthetic OCT corneal data is not available, synthetic corneas can be generated using the 3D corneal model to augment the four human subjects that were sampled in the University of Waterloo, Department of Physics. For each synthetic cornea, ground truth is randomly generated based on a statistical analysis of OCT data. Then a synthetic cornea is generated using the ground truth.

5. **Reconstruct cornea from tomograms without tomogram motion correction**

   In Chapter 7, for the first time, a reconstruction combining corneal anatomy and OCT imaging data can be obtained using the 3D corneal model and tomograms from

each human and synthetic subject. This first reconstruction can estimate unsampled cornea regions and produce tomograms at any cross-section within the reconstructed cornea using all of the available tomogram data for a particular subject. This reconstruction does not attempt to estimate the pose, position and orientation, of the tomograms based on eye motion. Consequently, this reconstruction process dose not distinguish between measurement uncertainty and uncertainty due to eye-motion.

6. **Reconstruct cornea from tomograms with tomogram motion correction**

    Chapter 8 expands the reconstruction process of Objective 5 to explicitly account for tomogram pose error due to eye-motion during scanning, which is not corrected using existing corneal reconstruction techniques. Through an optimization formulation, the pose of each tomogram is iteratively determined through the minimization of reconstruction error.

7. **Extract corneal layer thickness measurements**

    Ever since the 1990s, researchers have wanted to view the cornea from any arbitrary perspective [79]. With the proposed reconstruction method, corneal layer thickness measurements can now be acquired using anatomical constraints and eye-motion correction in Chapter 7 and Chapter 8. Unlike existing methods, the proposed reconstruction method also utilizes these measurements to quantify the accuracy of the reconstruction when these measurements are compared with ground truth, obtained through manual segmentation of human data and through the creation of synthetic data.

Document organization

This document is divided into five parts. The first part, Chapter 2, provides detailed background on statistical reconstruction theory and corneal reconstruction. The second part, Chapter 3, provides a detailed description of the reconstruction challenges and detailed objectives. The third part, Chapter 4, expands upon existing reconstruction methods introducing additional notation, non-linear model decoupling strategies, and tomogram pose estimation. The following three chapters, Chapter 5, Chapter 6, and Chapter 7, make up the fourth part and implement model-based reconstruction for corneal medical imaging. The fifth part, Chapter 8, reduces reconstruction error using an optimization formulation to better estimate pose parameters for corneal tomograms.

Chapter 2

Background

This chapter is designed to introduce notation and concepts that are used throughout this document. While some sections are fundamental to understanding, many of the sections are referred to in the following chapters and are designed to provide additional detail or motivation that may not be necessary to thoroughly understand on the first read through. It is recommended that the reader have a high level familiarity with the background material and to review specific sections and notation as they are referred to throughout the other chapters.

Chapter organization

This background is divided into two main sections that describe statistical reconstruction theory and corneal background.

The first section, Section 2.1, describes existing statistical reconstruction theory and the components involved in creating a reconstruction from measurements and a model of the subject. Since the entire document is based on this section, a good understanding of the concepts is recommended.

The second section, Section 2.2, provides background specific to reconstruction applied to corneal imaging. This section describes the anatomy of the cornea, a corneal imaging system, phsyological motion that occurs during imaging, existing reconstruction methods, and existing methods to extract useful measurements. This section provides supporting background for chapters 5 to 8. On the first read through, it is recommended that the reader be familiar with how the corneal structure appears after being imaged and that eye

motion during imaging can create measurement location uncertainty when attempting 3D reconstructions. Any other detailed background is referenced explicitly in later chapters.

2.1 Statistical reconstruction theory

The statistical model-based reconstruction theory is designed to recreate an arbitrary subject, such as a cornea, from a series of measurements. During the reconstruction process, the properties of the arbitrary subject are estimated such that measurements obtained from the reconstructed subject are consistent with measurements from the real subject. Subsection 2.1.4 discusses many existing techniques that incorporate probabilistic models and deterministic models.

The method presented here formulates a subject model explicitly decoupling non-linear relationships and simplifying the reconstruction process. Additional states and a parametric non-linear, or linear, decoupling function can separate the underlying subject states into sets of decoupled states. The decoupling function can then be integrated into the forward model, Subsection 2.1.3, allowing existing reconstruction processes to estimate the decoupled subject states. Chapter 7 applies this decoupling technique to reconstruct a cornea given a non-linear constrained corneal model. Further, an estimate of the forward model parameters can be generated through an optimization formulation, minimizing the reconstruction error, improving the reconstruction process, and providing the imaging operator with motion corrected parameters.

Section organization

This section presents the theory in recommended implementation order: obtain measurements, understand and model the subject, relate the measurements to the subject model, reconstruct the subject from the measurements, and improve estimates related to the measurement process using the reconstruction.

2.1.1 Measurements

The first step when attempting to reconstruct an object is to obtain measurements. Since many reconstruction operators utilize matrix multiplication, lexicographic reordering [42, 47] allows arbitrarily sized measurement fields to be converted into and from measurement vectors.

Measurements are denoted using the vector $\mathbf{m}$. It is useful to distinguish subsets of $\mathbf{m}$, which are denoted as $\mathbf{m}_i$,

$$\mathbf{m} = \begin{bmatrix} \mathbf{m}_0 & \mathbf{m}_1 & \mathbf{m}_2 & \cdots & \mathbf{m}_i & \cdots \end{bmatrix}^T. \tag{2.1}$$

Then element k within subset i can be refereed to as $m_{i,k}$,

$$\mathbf{m}_i = \begin{bmatrix} m_{i,0} & m_{i,1} & m_{i,2} & \cdots & m_{i,k} & \cdots \end{bmatrix}^T. \tag{2.2}$$

Then the vector $\mathbf{m}$ can be written in terms of all of its elements,

$$\mathbf{m} = \begin{bmatrix} m_{0,0} \\ m_{0,1} \\ m_{0,2} \\ \vdots \\ m_{1,0} \\ m_{1,1} \\ m_{1,2} \\ \vdots \\ m_{i,0} \\ m_{i,1} \\ m_{i,2} \\ \vdots \end{bmatrix} \sim \mathcal{N}(0, R), \tag{2.3}$$

and are assumed to be Gaussian distribution with a variance of R.

Using the notation defined for $\mathbf{m}$, an arbitrary 2D matrix containing measurement data, M, can be rewritten into a vector using lexicographic reordering [42, 47]. Each element (a, b) in M, $M_{a,b}$, can be reordered and stored in $\mathbf{m}$,

$$\mathbf{m} = \begin{bmatrix} M_{0,0} & M_{0,1} & M_{0,2} & \cdots & M_{1,0} & M_{1,1} & M_{1,2} & \cdots \end{bmatrix}^T. \tag{2.4}$$

Vector $\mathbf{m}$ can also be written using each row of M, $\mathbf{M}_{a,:}$,

$$\mathbf{m} = \begin{bmatrix} \mathbf{M}_{0,:}^T \\ \mathbf{M}_{1,:}^T \\ \mathbf{M}_{2,:}^T \\ \vdots \end{bmatrix}. \tag{2.5}$$

9

For notational purposes, lexicographical reordering can simply be denoted using $M_{:}$,

$$\mathbf{m} = M_{:}. \tag{2.6}$$

Similarly, an arbitrary 3D field, Z can be lexicographically reordered into a vector $\mathbf{z}$,

$$\mathbf{z} = Z_{:} \tag{2.7}$$

$$\mathbf{z} = \begin{bmatrix} Z_{0,0,0} \\ Z_{0,0,1} \\ Z_{0,0,2} \\ \vdots \\ Z_{0,1,0} \\ Z_{0,1,1} \\ Z_{0,1,2} \\ \vdots \\ Z_{1,0,0} \\ Z_{1,0,1} \\ Z_{1,0,2} \\ \vdots \\ Z_{1,1,0} \\ Z_{1,1,1} \\ Z_{1,1,2} \\ \vdots \end{bmatrix}. \tag{2.8}$$

This reordering is generalizable to any multi-dimensional field.

2.1.2 Subject model

A subject model describes the object that is being reconstructed. Without a subject model, there is nothing to reconstruct. In general this model can contain any number of parameters, or states, and a mathematical definition used to generate a representation from the states. For example, a deterministic subject model of a sphere may consist of parameters representing the radius, ρ, the center, $\mathbf{c}$, and a mathematical definition relating all of the points, $\mathbf{x}$, on the sphere,

$$\left(\rho^2 - \left(x_0 - c_0 \right)^2 - \left(x_1 - c_1 \right)^2 - \left(x_2 - c_2 \right)^2 \right) = 0. \tag{2.9}$$

While there is an infinite number of possible subject models, when designing a subject model it is important to understand how the model is used and what measurements are available such that the relationship between the measurements and the model parameters, or states, can be straightforward.

There are a large number of resources that describe the modelling process for various applications, such as ecological rodent populations [11], rainfall distribution [20], archeology [12], and large number describing general modelling methods [23, 38, 42].

For the application of corneal imaging, there are several useful tools that merit discussion and definition: Markov random fields and model state decoupling.

A random field is an ordered collection of random variables [27, 60, 72, 112]. The states for a subject model can be stored as random variables within the field and the relationships between the states can be defined mathematically in order to complete the model.

Discrete verses continuous

Random fields can be represented either continuously or discretely. This document focuses on spatially discrete random fields containing continuous states; a discrete number of continuous states are stored in a multi-dimensional matrix, Z. While a continuous random field can represent Z with infinite precision, a continuous random field has practical computing limitations when trying to estimate Z. A discrete random field has the advantage of storing every state in Z for use in calculations. The computational ease of working with a spatially quantized representation of a continuous field is preferred to working with a continuous field directly; there are a finite number of states that require estimation. The decision to move from continuous to discrete random fields reduces computational complexity at the cost of increased use of computer memory.

For convenience, the matrix Z can be lexicographically reordered into the vector $\mathbf{z}$. Each element in $\mathbf{z}$ represents a state of the subject model. The value of state j, z_j is application specific and, for example, in an image processing context, can contain pixel intensity values, representing values in a texture [135], satellite imaging data [126], brain magnetic resonance images [147], and retinal layer boundaries [63].

Markov random field

Using a Markov random field, the subject model can be decoupled so that probability distribution of one state, z_j, is conditional on the states in its neighbourhood, the states

surrounding z_j [42]. The neighbourhood can be extended to any size, but the relationship must be reciprocated, forcing the shape to be symmetric [42]. If the field is also stationary, the statistical relationships between neighbouring elements are the same at all locations within the field [60,72]. Thus once the statistics for a single neighbourhood are determined, the statistics for the entire field are known and the storage complexity of a problem can be greatly reduced.

A covariance matrix, P, relating all of the elements of a stationary Markov random field z_j, has a special property that the relationships contained in the first column of P encompass the relationships of the entire field. These relationships can be compactly written as the matrix kernel [42], $\mathcal{P}$,

$$\mathcal{P}_: = P_{:,1}. \tag{2.10}$$

The number of non-zero elements in $\mathcal{P}$ represents the order of the Markov random field.

The Markov random field subject model is defined by relationships $\mathcal{P}$ and the discrete states, Z.

Deterministic random field

Unlike a Markov random field based on statistical relationships, a deterministic random field relates states in terms of constraints. In a linear model, a local constraint matrix, L, relates each neighbouring element in Z [42]. While a hard constraint matrix would ensure that

$$L\mathbf{z} = 0, \tag{2.11}$$

requiring every local constraint to be satisfied, the soft constraints can be represented as a minimization criterion [42],

$$\mathbf{z}^T L^T L\mathbf{z}, \tag{2.12}$$

allowing deviations from the constraint when they conflict with measurements. Of course, if the real world and ideal model are identical, the minimum of (2.12) occurs when (2.11) is satisfied. The squared constraints, Q, are represented as a positive-semidefinite matrix and are defined as

$$Q = L^T L. \tag{2.13}$$

Like a stationary Markov random field, a stationary deterministic field can be represented using a matrix kernel. The matrix L has a corresponding kernel $\mathcal{L}$ [42], defined such that,

$$\left(\mathcal{L} * Z\right)_{:} = Lz_j. \tag{2.14}$$

The matrix Q can be represented as a kernel, $\mathcal{Q}$, using a similar relationship. A complete stationary deterministic field subject model can be represented using $\mathcal{L}$ and Z or by $\mathcal{Q}$ and Z.

Second-order deterministic model

The following briefly derives a 3D second-order deterministic model, $\mathcal{Q}$, by convolving many first-order constraints with each other. The deterministic random field constraint kernels $\mathcal{L}_0$, $\mathcal{L}_1$, and $\mathcal{L}_2$ are defined as first-order constraints in each respective dimension as

$$\mathcal{L}_{0,0,0,0} = -1, \qquad \mathcal{L}_{1,0,0,0} = -1, \qquad \mathcal{L}_{2,0,0,0} = -1, \tag{2.15}$$
$$\mathcal{L}_{0,1,0,0} = 1, \qquad \mathcal{L}_{1,0,1,0} = 1, \qquad \mathcal{L}_{2,0,0,1} = 1. \tag{2.16}$$

The second-order constraints are defined by convolution,

$$\mathcal{L}2_0 = \mathcal{L}_0 * \mathcal{L}_0 \tag{2.17}$$
$$\mathcal{L}2_1 = \mathcal{L}_1 * \mathcal{L}_1 \tag{2.18}$$
$$\mathcal{L}2_2 = \mathcal{L}_2 * \mathcal{L}_2, \tag{2.19}$$

and $\mathcal{Q}$ is defined as the sum of convolutions,

$$\mathcal{Q} = \mathcal{L}2_0 * \mathcal{L}2_0 + \mathcal{L}2_1 * \mathcal{L}2_1 + \mathcal{L}2_2 * \mathcal{L}2_2. \tag{2.20}$$

In 3D, the kernel becomes

$$\mathcal{Q}_{:,:,0} = \mathcal{Q}_{:,:,4} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.21}$$

$$\mathcal{Q}_{:,:,1} = \mathcal{Q}_{:,:,3} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -4 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \tag{2.22}$$

13

and

$$\mathcal{Q}_{:,:,2} = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -4 & 0 & 0 \\ 1 & -4 & 18 & -4 & 1 \\ 0 & 0 & -4 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}. \tag{2.23}$$

Model state decoupling

One of the largest challenges of modelling is dealing with coupled states. Coupling refers to the interaction between states within the random field. When estimating coupled states there are a number of dependencies, either statistical or deterministic that must be dealt with. For example, when solving a system of linear equations, $Ax = b$, the matrix A is generally dense. Assuming that a unique solution exists, the linear model, A, can be decoupled through a change of basis, using eigenvalue decomposition [131], $A = V\Lambda V^T$. Once decoupled, the solution to the linear system is straightforward since $V^{-1} = V^T$ and Λ is a diagonal matrix, $x = A^{-1}b = V\Lambda^{-1}V^Tb$. In this example, the states can also be solved as a sequentially coupled system [42] through Gauss-Jordan elimination [131], once one state is known, the next state can be determined until all states are determined.

Although linear deterministic systems are well understood, decoupling linear models, and especially nonlinear stochastic models, are not trivial in general. States can be sequentially, acyclic, cyclic, or multilayer coupled [42]. A sequential coupling is when each state, expect for an initial state, is either stochastically or deterministically dependent on the previous state. This form of coupling is fairly straightforward because each state is only dependent on one other state. An acyclic coupling occurs when each state is dependent on one or more previous states. The complexity increases as the number of dependencies increase since one state may have a probability conditioned on every other state. Acyclic coupling can be causal or anti-causal [42].

Cyclic coupling extends sequential coupling by allowing future states to be dependent on past states and past states to be dependent on future states [42]. Cyclic coupling is non-causal. Estimating parameters for cyclic coupling models is challenging since solving this type of problem can involve estimating one state, propagating its changes, and then estimating another state. Unless the states reach equilibrium, these changes can propagate indefinitely. Unlike cyclic coupling, multilayer coupling allows any state to be dependent on any other state and is usually defined through a hierarchy of dependencies. Trying to estimate a state that is dependent on every other state, and knowing that simply modifying

14

the state value can propagate to any other state, which in turn propagates back cyclically, makes solving multilayer coupling problems the most challenging [42].

Depending on the type of coupling, a model might be decoupled by modifying the basis or through restructuring their dependence on scale, time, or hierarchical organization [42]. The structure of each modelling application dictates the correct decoupling approach. Assuming that a decoupling function, ℓ exists, if Z is coupled, ℓ can separate $\mathbf{z}$ into decoupled states $\mathbf{z}_1$, $\mathbf{z}_2$, $\cdots$,

$$\begin{bmatrix} \mathbf{z}_1 \\ \mathbf{z}_2 \\ \vdots \end{bmatrix} = \ell(\mathbf{z}). \tag{2.24}$$

For corneal reconstruction, decoupling is further investigated in Section 4.2 and Chapter 6 to find functions of ℓ that can decouple a corneal model.

2.1.3 Forward model

This section describes general components of the forward model. The forward model should account for the subject structure, the subject behaviour, and subject imaging techniques. The application specific forward model requires intimate understanding of the subject matter and the imaging process; often the forward model is iteratively constructed, starting with a simple model and increasing the complexity at each iteration as additional subject knowledge is introduced into the model. However, once a forward model is created, the relationship between the subject and obtained data can be exploited to correct for measurement errors in the data, infer missing measurements, and create synthetic models using the knowledge gained of the subject structure and behaviour

The forward model, f, is a function that relates the measurements, $\mathbf{m}$, to the state variables, $\mathbf{z}$,

$$\mathbf{m} = f(\mathbf{z}). \tag{2.25}$$

In general, a forward model is an application specific arbitrary non-linear function. However, if the subject model can be designed with the goal of simplifying the measurement and state relationship, the forward model can be simplified. Ideally, if there is a direct correspondence from $\mathbf{m}$ to $\mathbf{z}$, the forward model can be simplified and reduced to a sparse correspondence matrix, C,

$$\mathbf{m} = C\mathbf{z}. \tag{2.26}$$

2.1.4 Reconstruction

Given a set of measurements, $\mathbf{m}$ and a forward model, f, reconstruction determines the most likely set of states, Z, that minimize, $\tilde{Z}$, the difference between the true states, Z, and the estimated states, $\hat{Z}$,

$$\tilde{Z} = Z - \hat{Z}. \tag{2.27}$$

Since the true value of Z is rarely known, the reconstruction process can minimize the residual error, $\tilde{\mathbf{m}}$, the difference between $\mathbf{m}$ and the measurement value corresponding to $\hat{Z}$, $\hat{\mathbf{m}}$,

$$\hat{\mathbf{m}} = f\left(\hat{Z}\right) \tag{2.28}$$

$$\tilde{\mathbf{m}} = \mathbf{m} - \hat{\mathbf{m}}. \tag{2.29}$$

It is also difficult to calculate Z directly from f and $\mathbf{m}$, since an explicit inverse model, f^{-1}, may not exist, may be highly sensitive to noise, or otherwise difficult to calculate.

Estimators

The Bayesian Estimator (BE) [42, 123] generates $\hat{Z}$, an estimate of an unknown field Z that minimizes the expectation of some cost function $C(Z, \hat{Z})$ given measurements $\mathbf{m}$,

$$\hat{Z} \triangleq \arg_{\hat{Z}} \min \mathbb{E}\left[C(Z, \hat{Z})|\mathbf{m}\right]. \tag{2.30}$$

When $C(Z, \hat{Z})$ is the quadratic cost function,

$$C(Z, \hat{Z}) \triangleq (Z - \hat{Z})^T (Z - \hat{Z}), \tag{2.31}$$

(2.30) becomes the Bayesian Least Squares Estimator (BLSE),

$$\hat{Z} \triangleq \arg_{\hat{Z}} \min \mathbb{E}\left[(Z - \hat{Z})^T (Z - \hat{Z})|\mathbf{m}\right]. \tag{2.32}$$

The Bayesian linear least-squares estimation [42, 123], makes the assumption that the relationship between Z and $\mathbf{m}$ is linear, further simplifying (2.32),

$$\hat{\mathbf{z}} = \left(C^T R^{-1} C + P^{-1}\right) C^T R^{-1} \left(\mathbf{m} - C\mu_{\mathbf{z}}\right) + \mu_{\mathbf{z}}. \tag{2.33}$$

2.2 Corneal background

This section provides background into corneal imaging by first introducing corneal anatomy and then by understanding the ultra-high resolution optical coherence tomography UHROCT imaging system. Since the reconstruction method requires data obtained from UHROCT, it is important to be familiar with the advantages and limitations of the imaging process; for instance, the benefits of a 3D reconstruction over a 2D reconstruction are robustness through the understanding of imaging noise. To build a structural model of the cornea and to segment important corneal regions, the anatomy of the human cornea must also be understood; the anatomy provides the basis for a mathematical constraints. Further, understanding physiology, in particular eye-motion during imaging, allows reconstruction to account for measurement location uncertainty during the imaging process. It is possible for the eye to slowly move during imaging causing the location of the last measurements to differ from their expected location in reference to the first measurements. Therefore the following subsections provide necessary background to justify modelling decisions made in the modelling and reconstruction sections that follow.

2.2.1 Anatomy

To facilitate sight, the human eye consists of several components integrated into a visual system, as illustrated in Fig. 2.1. The light enters the eye through the cornea and passes into the aqueous, a clear fluid which maintains a constant pressure within the eye. Then the iris changes the size of the pupil to regulate the amount of light that passes from the aqueous and into the lens. The lens can change shape to allow the eye to focus on both near and far objects. Once through the lens, the focused light passes through the vitreous, which fills the centre of the eye and primarily consists of water, and lands on the retina. The retina consists of photoreceptors that are stimulated by the light. The optical nerve is responsible for sending the photoreceptors signals to the brain for processing and interpretation [125].

The cornea is the transparent tissue covering the front surface of the eye serving multiple purposes: it provides nearly 70% of the eye's refractive power [94], protects the eye from foreign matter [93], can act as a minimally invasive pathway for therapeutic drug delivery [70] to the lens and retina, and protects the inner eye from over-exposure to sunlight [7]. The cornea is the first structure that light passes through when entering the eye and as such, defects in the corneal structure can prevent the incoming light from correctly interacting with the retina [55, 133, 141]. A reconstruction of the cornea can play an

Table 2.1: Distance measurements used to create Fig. 2.1 are referenced.

Distance Measurements in the Human Eye	
Measurement	Distance [mm]
fovea to lens	17 [15, 118]
lens major axis	10 [45]
lens minor axis	4 [45]
anterior pole to retinal Bruch's membrane	24 [83]
anterior pole to center of rotation	13.5 [83]
corneal posterior curvature	6.5 [142]
corneal anterior curvature	7.8 [142]

important role in by allowing researchers to better identify these structural defects. The following discussion focuses on the corneal structure, the light's optical flow through the cornea, and discussion of disease and disorder associated with abnormal corneal structure.

Structure

Although the cornea consists mostly of water and protein, there are six distinct layers that can affect the path of light through the cornea, the epithelium, Bowman's layer, the stroma, Dua's membrane, Descemet's membrane, and the endothelium layer, illustrated in the rightmost panel of Fig. 2.2. Each distinct corneal layer is composed of different cells and collagen fibers that are stacked from the outer to the inner surface. The discovery of Dua's membrane was only published in September, 2013, and has yet to be confirmed by other researchers [34].

Corneal layers

The thickness for each layer is presented in Table 2.2. The epithelium layer consists of six layers of cells, the outermost cells constantly replaced by quickly regenerated interior cells [130]. The epithelium has the highest refractive index in the cornea and separates air from the inner layers [14]. Bowman's layer consists of tough collagen fibers that protect the stroma [6]. The stroma, the largest corneal layer, consists of collagen fibers and keratocytes. While the collagen fibers provide structure for the stroma, the keratocytes repair damage to the stroma [145]. The Descemet's membrane consists of collagen fibers that are less rigid than those found in the stroma and Bowman's layer [6]. The Descemet's membrane

(a)
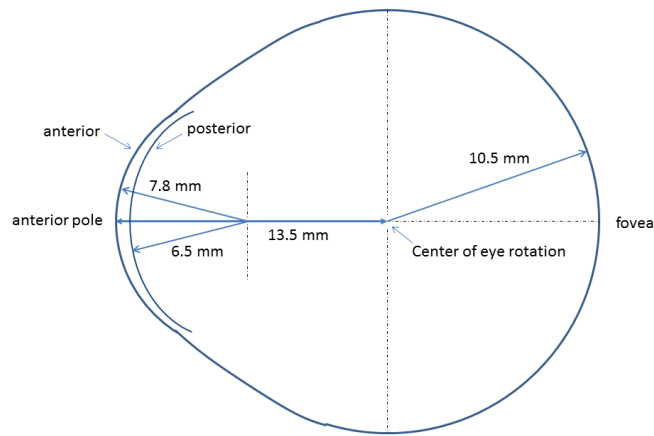
(b)

Figure 2.1: The human eye consists of many components that affect the passage of light. The major components of the eye are labelled above (a), the cornea, the iris, the pupil, the lens, and the retina. The anterior pole and fovea are also labelled for reference. Additional distances are labelled in (b) illustrate the anterior and posterior corneal curvature near the anterior pole and the distances from the anterior pole and fovea to the eye's center of rotation. All of the distances used to construct this diagram are referenced in Table 2.1.
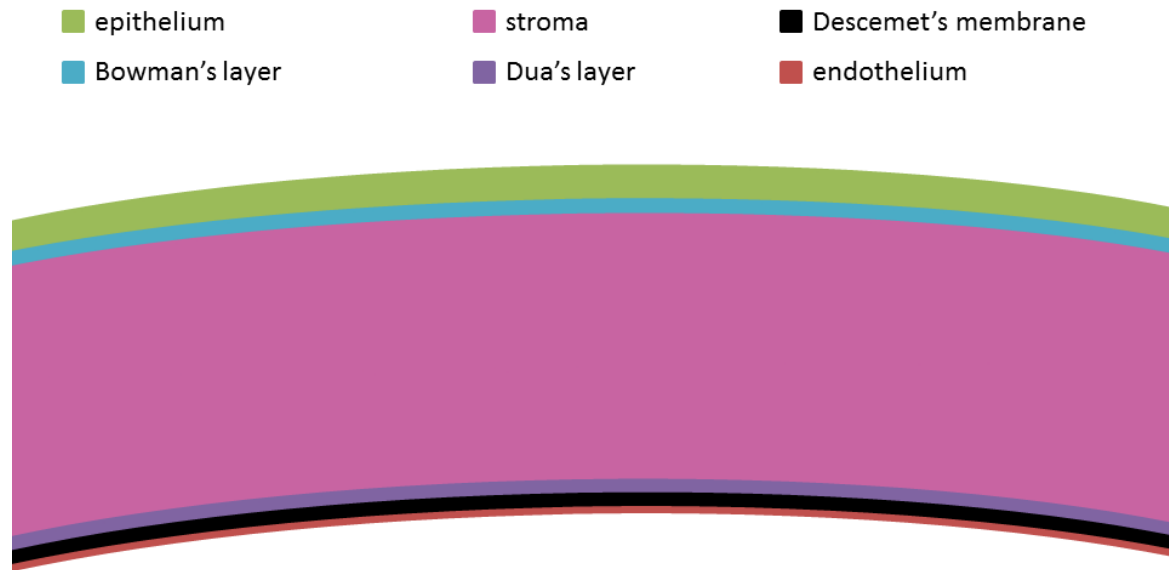
Figure 2.2: A cornea's structure consists of six layers: the epithelium, Bowman's membrane, stroma, Dua's layer, Descemet's membrane, and the endothelium.

allows fluid to pass into the stroma for hydration and prevents stroma swelling, evacuating fluid using "ionic pumps" [95]. The collagen in the Descemet's membrane is produced by the hexagonally shaped endothelium cells [111]. Unlike the other corneal layers, the endothelium cells do not regenerate; instead when endothelium cells die, the surrounding living cells stretch to cover any gaps [130].

Transparency

A transparent and functional cornea is highly dependent on the shape and structure of the layers. Even though the stroma is made up of non-transparent cells, representing about 90% of the cornea thickness [69], there are two prominent hypotheses explaining how the cornea maintains overall transparency. Maurice [81] describes transparency through destructive interference relating to the cellular lattice structure of the cornea while Goldman [50] expands on Maurice showing that the spacing of collagen fibers must be less than $200nm$ for transparency to occur. Both theories indicate the importance of structure required for a transparent cornea.

Table 2.2: The mean central thickness in μm of each corneal layer for an adult is presented below. UHROCT can resolve 5 μm in the axial direction, making it difficult to resolve the endothelium layer boundaries.

Layer	Thickness (μm)
epithelium	50 - 59.1 [111, 117]
Bowman's layer	8 - 21.2 [111, 117]
stroma	424 - 500 [108, 111]
Dua's layer	6.3 - 15.8 [34]
Descemet's membrane	7.8 - 14.0 [34, 111]
endothelium	4 - 6 [106, 111]

Disease and disorder

A healthy cornea is designed to focus light through the iris and onto the lens. However, corneal disease and disorder can interfere by changing the shape and transparency of the cornea. Many of these diseases can be screened, treated, or avoided through the use of optical coherence tomography.

Keratitis is an inflammation of the cornea caused by bacterial, fungal, viral, or other foreign organisms [77]. The severity of inflammation may be superficial, only causing the epithelium to swell or may be more severe causing swelling throughout the entire cornea [107]. Corneal thickness measurements were historically used to study the amount of inflammation induced by keratitis [25] and, more recently, optical coherence tomography corneal thickness measurements were utilized to study reduction of inflammation due to treatment of bacterial keratitis [62].

Corneal abrasion is a loss of the epithelium layer, often in the form of a scratch, resulting from particulates, chemical burns, or a foreign object coming into contact with the cornea [139]. Since deep scratches can result in infection [139], there is great incentive to use non-contact imaging, such as optical coherence tomography, to diagnose and treat other diseases to avoid the risk of corneal abrasion [102].

While glaucoma is a disease that affects the optical nerve, not the cornea, glaucoma increases intraocular eye pressure [51]. Measuring corneal layer thickness provides a way to screen patients for risk; patients with thin corneal thickness are more sensitive to pressure than patients with thicker corneal layers [26, 51, 82]. An automated method to measure corneal layer thickness can provide medical personal with fast and accurate feedback allowing a technician to verify the procedure instead of manually measuring the cornea. Further,

since patients with thinner corneal thickness have higher risk of glaucoma independent of intraocular eye pressure and age, recent research is attempting to find the genes responsible for glaucoma by identifying genes responsible for thinner corneal layer thickness [137]. This research area can benefit from an automated layer measurements.

Corneal hypoxia occurs when the cornea is deprived of oxygen and is associated with a disorder called neovascularisation, causing blood vessels to grow into the cornea [109,145]. Since hypoxia is common in patients wearing contact lenses [145] lenses were designed to reduce the effect of hypoxia and discomfort on patients [132]. Hypoxia studies relay on corneal layer thickness measurements to determine how much corneal swelling occurs due to contact lens wear [10].

Keratoconus is a disorder in which the shape of the cornea resembles a cone [124]. Patients suffering from keratoconus complain of blurred vision and often require contact lenses [145]. The use of optical coherence tomography has also been used to diagnose moderate and severe keratoconus with the same accuracy as previous diagnostic tools, but has a significant advantage of being able to detect early onsets of the disorder due to the ability of optical coherence tomography to image corneal thickness [73]. Having a 3D reconstruction of the cornea, would allow keratoconus measurements to be obtained from the same location from the reconstruction reducing measurement location error due to eye-motion during scanning.

Fuchs' dystrophy causes the Descemet's membrane layer to thicken gradually over time causing blurred morning vision that worsens with age [92]. The increased resolution available from ultra-high resolution optical coherence tomography allows early screening of Fuchs' dystrophy since the thickness of the Descemet's membrane can be twice as thick in patients with the disease than healthy patients [121].

Measurements

Corneal packymetry is the methodology of measuring corneal layer thickness. From the discussion of corneal disorders and diseases corneal thickness measurements provide an important role in diagnostics and treatment. The accuracy of these measurements rely on consistently measuring the same location within the cornea. Since obtaining a 2D tomogram at a reproducible location within the cornea is difficult due to eye-motion [97], a 3D reconstruction can help determine the location of the tomogram allowing metrics to be measured from more consistent positions.

Layer thickness is measured as the distance between each layer boundary, Fig. 2.3. Corneal thickness is the distance between each corneal layer boundary measured in the
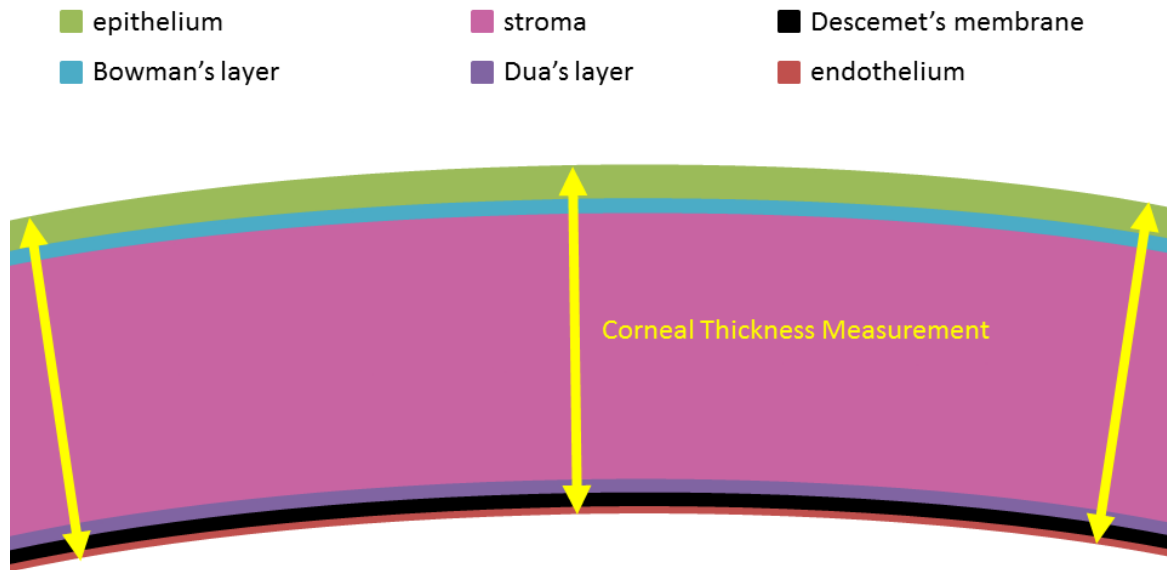
Figure 2.3: Layer thickness is measured as the distance between layer boundaries along the corneal axial direction.

axial direction [68,74]. It is also difficult to know exactly where in the cornea the tomogram is imaged; for instance it is difficult to know how far the imaging plane is placed relative to the apex of the cornea. Since the apex can be identified as the area in the cornea with high reflectivity, which correspond to imaging artifacts in the tomogram, Fig. 2.4. As a result, researchers rely on repeated measurements to estimate the central corneal thickness, usually in reference to the apex, in order to obtain reproducible measurements [91]; measuring layer thickness in a location other than the apex is unrepeatable, without registration. Once a 3D model is created for a cornea, corneal layer measurements can be consistently obtained at any position within the model.

Li applied confocal microscopy through focusing to measure the central layer thickness of the epithelium, Bowman's layer, and total corneal thickness [71]. Although the approach is limited to manually measuring the central thickness, it was the first technique to obtain measurements for three of the five corneal layers. In more recent work, Li was able to automatically measure the central corneal thickness using optical coherence tomography and the average of the prominent epithelium and endothelium boundary layer locations [74]. Unfortunately, this method is limited to only the central corneal thickness.

(a) reproduced from Baikoff [4]

(b)

(c)

Figure 2.4: The appearance of the cornea and corneal layer boundaries within OCT tomograms is presented to illustrate the presence of corneal features depending on scale. (a) A view at this scale contains features, such as the iris structure, that can be useful for reconstruction. (b) As the imaging magnification increases, there are few features remaining that can be used for registration. The corneal structure remains, but has similar appearance from many different perspectives. (c) As the imaging system nears the anterior pole, or corneal apex, high levels of scatter potential, indicated by pixels with low intensitiy (dark) appear. (a) is reproduced from Baikoff [4] and (b) and (c) are obtained from the University of Waterloo, Department of Physics.

Corneal reconstruction motivation

An accurate reconstruction of the cornea can assist in the study of corneal deformation. Diseases and disorders specific to the cornea may be detected or studied using a 3D model obtained from OCT images. By comparing current corneal layer thickness measurements to previously obtained values in the same corneal region, the progression of the disease or treatment can be measured. Further reconstruction can reduce eye-motion during imaging that would otherwise require specialized sampling schemes, such as repeated sampling of the same volume [97].

2.2.2 Ultra-high resolution optical coherence tomography

From the discussion of corneal diseases and disorders in Subsection 2.2.1 the usage of optical coherence tomography (OCT), an imaging technology, in diagnostics, treatment, and research is prominent [10,62,73,102,121,137]. OCT allows non-contact *in vivo* corneal imaging, which reduces risk of corneal abrasion and infection. Further, the evolution of corneal ultra-high resolution optical coherence tomography (UHROCT) allowed the cornea to be imaged with 6 μm axial resolution in 2005 [18] and 3.4 μm axial resolution in 2007 [19]. The data used throughout this study has 3.2 μm axial by 10 μm lateral resolution [57]. This resolution is a major improvement over regular OCT as it allows five individual corneal layers to be resolved [19], the epithelium, Bowman's layer, stroma, Descemet's membrane, and endothelium.

Appearance of cornea within UHROCT

Fig. 2.4 (a) shows an example of OCT imaging and Fig. 2.4 (b) and Fig. 2.4 (c) show the magnification capabilities of UHROCT imaging. The high levels of scattering potential in Fig. 2.4 (c) are one of the few reference points useful for reconstruction. The lack of prominent differences between Fig. 2.4 (b) and Fig. 2.4 (c) make it difficult to determine precisely where (b) is located relative to (c). Fig. 2.5 shows the location of each visible corneal layer within the UHROCT tomogram. Dua's layer is not labelled since expert data for an UHROCT tomogram is not yet available.

An UHROCT for a human cornea is represented in Fig. 2.5. Notice that the outer layer boundaries are clearly distinguishable from the surrounding fluid, the inner layers are more difficult to visualize because they have a low refractive index. Since inner layer boundaries have only suitable visual contrast with the presence of noisy cellular scattering potential,

819.2 μm

epithelium

Bowman's layer

stroma

Descemet's membrane

endothelium

10 mm
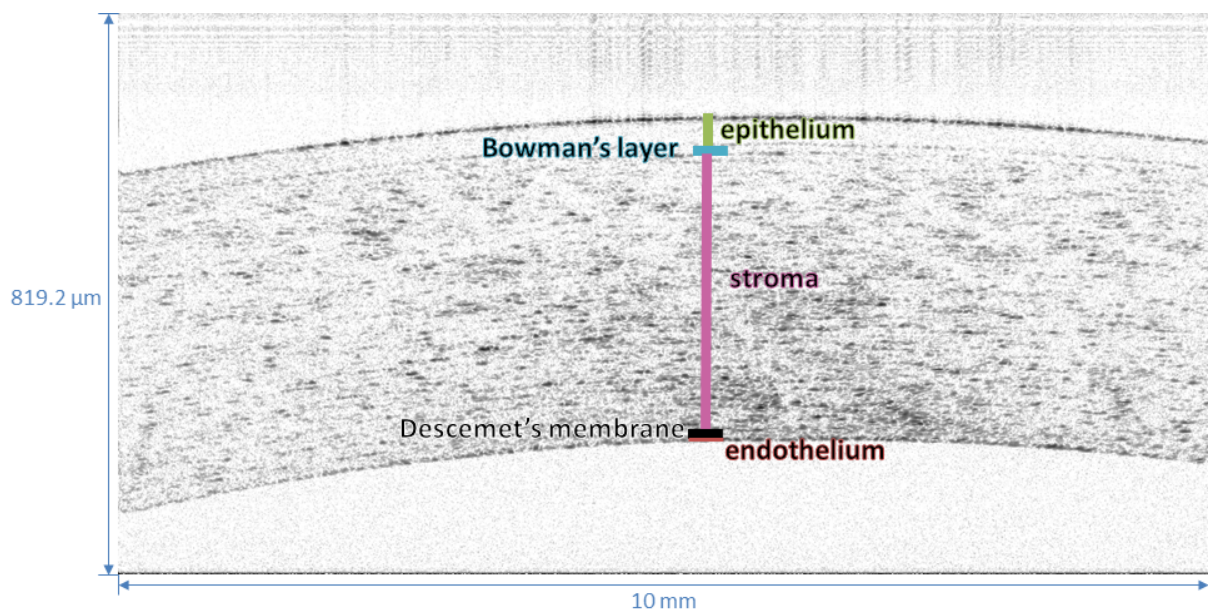
Figure 2.5: Five of the six corneal layers are visible and labelled in the UHROCT tomogram above. This tomogram has 3.2 μm axial resolution and 10 μm lateral resolution. The corneal layer dimensions are to scale.

locating inner layer boundaries is a challenge. Fig. 2.5 labels and highlights each visible corneal layer. Notice that while the outer corneal layers are clearly visible and can even be located with basic image processing filters [68], the inner corneal layers are much harder to segment.

UHROCT sampling time

It is important to profile the length of time required to capture a tomogram so that the amount of patient movement can be statistically modelled. When the UHROCT system samples a human cornea, a tomogram is constructed from a series of A-scans that sample the cornea in the axial direction. Each A-scan is represented as a vector containing intensities calculated from a narrow conical beam. A single image, a tomogram, produced from UHROCT imaging is referred to as a B-scan. The B-scan is a 2D matrix containing pixel intensities that are populated by two or more A-scans. Common B-scan format consists of 500 to 1,000 A-scans. Typical UHROCT imaging speeds can capture 29,000 to 47,000 A-scans per second [18,86]. A B-scan consisting of 1000 A-scans would require about 21 ms to 34 ms to generate a single tomogram, allowing 29 to 47 tomograms to be imaged per second.

OCT notation

This notation is essential since it is used throughout this document. A tomogram is a matrix, often stored as a bitmap, containing measurements of scattering potential [40], also referred to as scattering coefficient [144], and in some OCT imaging systems, equivalent to backscatter [116]. During an imaging session corneal thickness measurements are determined from a set of tomograms, T, for a single patient. The index variable i is used to denote an individual tomogram, T_i, from the set. Since bitamps are commonly used to store T_i, let $\eta_{s,0,i}$ and $\eta_{s,1,i}$ represent the number of pixels in the horizontal and vertical direction, respectively, within tomogram i. Since the dimensions of each pixel are dependent on the imaging system, let $\gamma_{s,0,i}$ and $\gamma_{s,1,i}$ denote the width and height, respectively, of T_i in μm. Using an UHROCT system with 3.2 μm axial and 10 μm lateral resolution, a tomogram with dimensions $(\eta_{s,0,i}, \eta_{s,1,i}) = (1000\ pixels, 256\ pixels)$ would sample a corneal cross-section with a size of $(\gamma_{s,0,i}, \gamma_{s,1,i}) = (10,000\ \mu m, 819.2\ \mu m)$. For future notational convince, let

$$\gamma_{s,i} = \begin{bmatrix} \gamma_{s,0,i} & \gamma_{s,1,i} \end{bmatrix}^T. \tag{2.34}$$

A position in a tomogram can be denoted using the vector s, which contains the position along the horizontal (s_0) and vertical (s_1) axis of a tomogram,

$$s = \begin{bmatrix} s_0 & s_1 \end{bmatrix}^T, \tag{2.35}$$

as illustrated in Fig. 2.6. Thus, an arbitrary position within T_i can be denoted as $\mathbf{s}_i$,

$$\mathbf{s}_i = \begin{bmatrix} s_{0,i} & s_{1,i} \end{bmatrix}^T, \tag{2.36}$$

and the position of the aribtrary pixel k, is denoted as $\mathbf{s}_{i,k}$,

$$\mathbf{s}_{i,k} = \begin{bmatrix} s_{0,i,k} & s_{1,i,k} \end{bmatrix}^T. \tag{2.37}$$

OCT principles

An UHROCT sampling model, that will only be used in this subsection, can be defined based on the physics of OCT. This discussion is included to supplement understanding of corneal imaging and is not necessary to understand other sections of this document. Further, the work by Fercher published in 2003 is an excellent resource if additional details are desired [40]. Since the data for this document was collected using low-coherence inter-ferometry (LCI) and Fourier-domain OCT (FDOCT), the discussion is limited to FDOCT implemented with LCI. See Fercher [40] for a discussion of other OCT system variants. All of the following notation is used exclusively for the remainder of this section.

All FDOCT require a coherent light source, Fig. 2.7 (a). For optimal resolution, the light source should have low temporal coherence, giving the system greater depth resolu-tion, and high spatial coherence, giving the system greater lateral resolution [40]. Let $A^{(i)}$ represent the amplitude of the coherent light source, $\phi(t)$ represent the phase at time t, λ represent the wavelength, ω represents the angular frequency, $\mathbf{k}^{(i)}$ represent the wave vector, and k represent the wave number defined as

$$k = \mid \mathbf{k}^{(i)} \mid = \frac{2\pi}{\lambda}. \tag{2.38}$$

Then the wave, $V^{(i)}$ can be described in terms of time t and position $\mathbf{r}$,

$$V^{(i)}(\mathbf{r}, \mathbf{k}^{(i)}, t) = A^{(i)} \exp\left(j\mathbf{k}^{(i)} \cdot \mathbf{r} - j\omega t\right), \tag{2.39}$$

by definition [40, 136], where $j = \sqrt{-1}$.

Figure 2.6: $T_i(\mathbf{s})$ is a point within the tomogram T_i. The physical width $\gamma_{s,0,i}$ and height $\gamma_{s,1,i}$ correspond to $\eta_{s,0,i}$ and $\eta_{s,1,i}$ pixels in the horizontal and vertical directions respectively.

Figure 2.7: A Fourier-domain optical coherence tomography (FDOCT) system contains (a) a coherence source that splits at (d) the fiber coupler and routed to (b) the sample and to the reference mirror. The reference beam reflects off of the reference mirror and returns to (e) before entering (d). The sample at (b) causes backscattering that returns alone the probe beam path to (c) before entering (d). The reference and probe beams are combined at (d) creating interference patterns and are routed to (f) the beam detector that measures the spectrum and sends the samples to a PC.

The scattering potential of the sample, $F_s(\mathbf{r}, k)$, at position $\mathbf{r}$ for the specified k, illustrated in Fig. 2.7 (b), is defined in terms of the sample's complex refractive index distribution, m, as

$$F_s(\mathbf{r}, k) = k^2 \big(m(\mathbf{r}, k)^2 - 1\big). \tag{2.40}$$

The scattered wave, V_s at distance d from the sample, resulting from $V^{(i)}(\mathbf{r}, \mathbf{k}^{(i)}, t)$ can be derived using the scattering vector $k^{(s)}$ and K, the difference between of $k^{(i)}$ from $k^{(s)}$ and the far field approximation to determine the amplitude of A_S, the scattered sampled wave [40],

$$V_s(\mathbf{r}, \mathbf{K}, t) = A_S \exp\big(j\mathbf{k}^{(s)} \cdot \mathbf{r} - j\omega t\big) \tag{2.41}$$

$$A_S(\mathbf{r}, \mathbf{K}, t) \approx \frac{A^{(i)}}{4\pi d} \int_{Vol(r')} F_s(\mathbf{r}') \cdot \exp(-jK \cdot \mathbf{r}') d^3 \mathbf{r}'. \tag{2.42}$$

The scattering vector $\mathbf{k}^{(s)}$ is the wave vector from the sample to the detector. Using a narrow beam, measuring backscatter, and assuming that F_s is independent in lateral

directions simplifies (2.41) and (2.42). The amplitude a_S and phase ϕ_S of the backscatter wave at location Fig. 2.7 (d) allows the wave to be written as [40]

$$V_S(z, K, t) = A_S(2k) \exp(-jkz - j\omega t) \tag{2.43}$$
$$A_S(2k) \approx a_s(2k) \exp\big(j\phi_S(2k)\big). \tag{2.44}$$

Since the sample scattering wave, A_S, can be shown to be proportional to the inverse Fourier transform of the scattering potential using (2.42) and by fixing the value of d [40],

$$A_S(2k) \propto \mathcal{F}^{-1}\big(F_s(z)\big) \tag{2.45}$$
$$F_s(z) \propto \mathcal{F}\Big(a_s(2k) \exp\big(j\phi_S(2k)\big)\Big). \tag{2.46}$$

For the particular FDOCT LCI imaging system used to collect data for this application, the scattering potential, $F_s(z)$, is proportional to the backscattering sample wave, $A_S(2k)$.

Let $|\gamma_{SR}|$ represent the degree of coherence between A_S and R, δ_{SR} represent the phase delay, I_S represent the spectral intensity of A_S, I_R represent the spectral intensity of R, ϕ_S represent the phase angle of A_S, and ϕ_R represent the phase angle of R. The interferogram, is created at Fig. 2.7 (d) where a reference beam, R, Fig. 2.7 (e), is combined with A_S, creating a function of the cross-correlation between A_S and R, G_{SR} [40],

$$G_{SR}(2k) = 2\big(I_S(2k)I_R(2k)\big)^{.5} |\gamma_{SR}| \Re\left(\exp\Big(j\big(\phi_S(2k) - \phi_R(2k)\big)\Big)\right). \tag{2.47}$$

The spectral interferogram, $I(2k)$ at the beam detector, Fig. 2.7 (f) is the result of combining A_S, R, and G_{SR} [40],

$$I_{SR}(2k) = I_S(2k) + I_R(2k) + G_{SR}(2k). \tag{2.48}$$

Using these relationships, Fercher shows that $I_{SR}(2k)$ can be found in terms of the z_R, the distance of the reference mirror to the beam splitter [40],

$$I_{SR}(2k) = |\hat{F}_s(2k)|^2 + R + 2\sqrt{R}|\hat{F}_s(2k)| \cos(2kz_R). \tag{2.49}$$

The distance $\mathbf{z}_R$ should be twice the desired sampling distance to avoid sample overlap [40]. The constant terms $|\hat{F}_s(2k)|^2$ and R can be removed from $I_{SR}(2k)$ by taking two measurements of the same location, ensuring that the second measurement is out of phase with the first, and subtracting the results [115], or by simply removing the DC gain from a spectrometer,

$$I_{SR}(2k) = 2\sqrt{R}|\hat{F}_s(2k)| \cos(2kz_R). \tag{2.50}$$

A measurement of $F_s(2k)$ can be determined from measurements and (2.50).

31

Table 2.3: The corneal eye motion resulting from tremors, drifts, and microsaccades is summarized. Unlike drifts and microsaccades, the tremors are aperiodic and wave-like, negating the tremor effect within each A-scan every $11ms$, approximately. The microsaccades have a larger range since specific patient and laboratory conditions can affect the effect of microsaccades, see Martinez et al [78] for details.

	units	epithelium	endothelium
tremor (amplitude)	μm	5.14	4.92
drift	$\frac{\mu m}{s}$	96.6	92.5
microsaccades (trained)	$\frac{\mu m}{s}$	$707 - 3299$	$676 - 3157$
microsaccades (typical)	$\frac{\mu m}{s}$	$908 - 4240$	$869 - 4060$
microsaccades (max)	$\frac{\mu m}{s}$	$6600 - 11500$	$6310 - 11000$

2.2.3 Eye motion

It is important to compensate for the eye motion while the human subjects are staring on a light in the UHROCT scanner during the UHROCT sampling process. This subsection describes three types of eye motion that occur while the patient is fixating at an object. Section 6.3 develops an eye motion model based on the contents of this subsection. Table 2.3 summarizes the range of motion that is detailed in the following discussion.

When focusing on an object, first the eye suddenly moves at a very high speed and then enters a relatively stable state for about 200 to $600ms$ [58] while the eye focuses on the target. While studying the effects of neural and visual fixation, Martinez-Conde details three primary eye-movements that occur during fixation, tremor, drifts and microsaccades [78]. In humans, the wavelike tremors occur at approximately $90Hz$ with an amplitude of approximately $1 - 4$ μm, the diameter of a cone in the fovia [16, 105, 146]. This amplitude measured at the retina can be converted into corneal amplitude at the anterior pole using the rotational center of the eye about a point approximately 13.5 mm behind the anterior pole (cornea) and the distance of 24 mm between the anterior pole and Bruch's membrane (retina) for a healthy adult [83],

$$r_{retina} = 24 - 13.5 = 10.5 \ mm \tag{2.51}$$

$$\beta_{retina} = \frac{1 \ \mu m}{10500 \ \mu m} = 9.52 \times 10^{-5} rad \tag{2.52}$$

$$\beta_{cornia} = \beta_{retina} = 9.52 \times 10^{-5} rad \tag{2.53}$$

$$r_{cornia} = 13.5 \ mm \tag{2.54}$$

$$d_{cornia} = r_{cornia}\beta_{cornia} = 1.29 \ \mu m, \tag{2.55}$$

resulting in a distance between 1.29 and 5.14 μm at outer corneal epithelium layer boundary. Similarly using the same derivation assuming an average central corneal thickness of 581 μm [8], the amplitude of the tremor is approximately 1.23 to 4.92 μm. Since the resolution of a corneal UHROCT of approximately 3×10 μm (axial $\times$ lateral) [10] the resulting tremor motion will account for any measurement to be offset by no more than 2 pixels in the axial direction and no more than 0.5 pixels in the lateral direction at any time.

A drift is a slow trending motion of the eye that, for the duration of the UHROCT sampling procedure, last about $0.2 - 1$ seconds [110] with an average speed of 24.6min per second [128] measured at the retina. This motion can be converted into corneal motion,

$$\omega_{retina} = 24.6\frac{min}{s} \tag{2.56}$$

$$\omega_{cornia} = \omega_{retina} = 24.6\frac{min}{s} = 7.16 \times 10^{-3}\frac{rad}{s} \tag{2.57}$$

$$r_{cornia} = 13.5 \; mm \tag{2.58}$$

$$s_{cornia} = r_{cornia}\omega_{cornia} = 96.6\frac{\mu m}{s}. \tag{2.59}$$

The drift motion causes the epithelium of the cornea to move at approximately 96.6 μm per second. Using the resolution of a corneal UHROCT on average the eye can drift in the lateral direction by 9.6 pixels per second near the epithelium layer and, using the same derivation assuming an average central corneal thickness of 581 μm [8], a drift of 92.5 μm per second or 9.2 pixels per second near the endothelium layer.

Although the role of microsaccades is contested [78], this motion is the most significant of the three eye motions and takes the form of short and quick eye motion. While it is possible to train patients to suppress microsaccades under some circumstances, a typical human experiences microsaccades at a frequency between 0.23 Hz to 2.5 Hz [78] with an average speed of 3 to 14 deg per second [31, 146, 149]. Again, this retinal measurement is converted into a corneal measurement resulting in 707 to 3299 μm per second and 676 to 3157 μm per second at the epithelium and endothelium boundaries respectively. Although microsaccades occur less frequently than drifts, when microsaccades occur they can account for a retinal speed of 706 to 3300 μm per second [87], or 908 to 4240 μm per second and 869 to 4060 μm per second at the epithelium and endothelium boundaries respectively. Regardless that some theories suggest that microsaccades correct for drift and that microsaccades return the subject of focus to the center of view [78], a max speed of 28 to 49 deg per second, or measured at the cornea, 6600 to 11500 μm per second, or 6310 to 11000 μm per second at the epithelium and endothelium boundaries respectively.

Throughout the entire tomogram, a trained patient results in about 71 to 330 pixels per second and an average patient results in 91 to 424 pixels per second in the lateral direction. The axial motion is 236 to 1100 pixels per second and 303 to 1410 pixels per second for a trained and average patient, respectively.

The maximum microsaccade motion is significant, 1155 pixels per second laterally and over 3800 in the axial direction. It is assumed that the UHROCT operator discards tomograms containing such large obvious motions, when a corneal layer boundary is visible at the start of sampling but not at the completion of sampling, for example. However, the compounded maximum microsaccade motion may be less obvious after several sequential scans.

Table 2.3 summarizes the eye motion resulting at the outer corneal epithelium and endothelium layer boundaries. The drift and microsaccade motion is illustrated in Fig. 2.8, reproduced from Pritchard [100]. Fig. 6.22 illustrates how the human eye rotates about all three axes [9]. The rotational eye movement causes each point on the corneal layer boundary to translate in all three directions with respect to the UHROCT imaging system.

2.2.4 Reconstruction methods

Corneal reconstruction is the process of converting sets of corneal data into a 3D model. Unfortunately existing reconstruction techniques attempt to align sequential sets of images by pixel correspondence and then proceed to stack the images without any regard for eye-motion during scanning and little regard to estimating unsampled corneal regions. In corneal literature, reconstruction implies reshaping corneal measurements into a 3D data structure requiring each voxels in the volume to be measured directly.

Corneal reconstruction

In 1990, Masters constructed an optical dataset of a rabbit cornea using confocal microscopy and sampled every voxel of a $256 \times 256 \times 134$ volume with 3 μm resolution in each dimension [79]. The first 3D reconstruction and visualization of a rabbit cornea was created directly from the confocal microscopy data. The authors created the reconstruction to allow other researchers to view their dataset from multiple perspectives and the authors were able to take arbitrary cross-sections of their data within the volume. This was a major feat in 1990. Like Masters, the purpose of reconstruction in this document is also to allow arbitrary cross-sections and views of corneal data. However instead of relying on measurements for each voxel, the definition of reconstruction throughout this

Figure 2.8: The drift and microsaccade motion are illustrated on the retina. The diameter of the fovea above is 50 μm. The jagged curves results from the drift while the lines represent the microsaccade motion. This figure is reproduced from Pritchard [100].

Figure 2.9: The human eye rotates on all three axes [9] about a point 13.5 mm [83] behind the anterior pole. The curvature of the cornea with respect to the curvature of the eye is not to scale and exaggerated to illustrate that eye rotations cause the distance of the UHROCT imaging system to the cornea, and all the corneal layers, to vary.

document also utilizes corneal anatomy and physiology to estimate missing data, reducing noisy measurements, and accommodating movement during the imaging process.

In 2002, Vabre created a 3D reconstruction of a tadpole using an optical coherence tomography system [134]. The process required 300 tomograms to sample a volume that was 360 $\mu m \times$ 360 $\mu m \times$ 200 μm achieving an axial resolution of 1 μm. Like Masters, this process was reconstructed by populating all of the voxels with tomogram data. Although the authors wrote their own visualization software, there was no use of a corneal model or attempt to reduce measurement error during reconstruction. The relationship between Z and $\mathbf{m}$ is one-to-one.

In 2007, Scarpa presents a method to reconstruct a human cornea from confocal microscope imaging [114] with the assumption that each corneal image can be translated in the x-y plane. A region of interest is identified in each sequential set of images, then a normalized correlation method [28] is applied to the region of interest to find correspondences between the image frames. The images in the stack are translated to align the correspondences in consecutive images. The process, however, does not use corneal antaomy for reference, and does not account for eye-motion, only shifts in the x-y plane. The method also constrains the sampling process, requiring densely sampled parallel sets of medical images to be obtained. The process assumes that each layer is separated by a single pixel, and fills in missing layers using linear interpolation between the two nearest images. Their reconstruction process does not take into account measurement noise and without a corneal model, does not provide any justification to their interpolation scheme.

In 2009, Zhivov presented a reconstruction method from a confocal microscope using a software package AMIRA 3.1 [148]. Like Scarpa, the image stacks were aligned. Zhivov used a semi-automatic least-square algorithm to align the stack of images before reconstruction. The volumetric renderings of the 3D cornea using AMIRA have a nice appearance, but the authors should be cautious using such reconstructions. There is no discussion of how missing measurements are determined, no discussion of any underlying model used for reconstruction, and no discussion of reconstruction error. It is very likely that the reconstruction error is in fact zero indicating that the reconstruction corresponds exactly to perfect measurements and that eye motion during imaging was not considered. A proper reconstruction should have residual error because measurements are imprecise and the lack of discussion from Zhivov indicates that measurement error has been overshadowed by nice visualizations.

AMIRA is a popular image stacking tool used to generate 3D views of the cornea. The software provides a suit of tools that can be used to align a stack of OCT images by comparing the direct image intensity and any salient features contained in sequen-

tial images [106, 129]. The software also allows the users to manually align the images. Unfortunately, the package does not use the structural properties of the cornea in the reconstruction process, preventing a suitable 3D reconstruction, and the generated 3D reconstructions failed to yield the accuracy necessary for corneal layer thickness research.

While pixel intensity is a prominent feature in corneal images, it is not unique and should not be used in isolation without other techniques.

Non-corneal reconstruction

There are several examples of reconstruction techniques in other medical fields that develop models to account for measurement uncertainty and to estimate unsampled regions of the subject. Some of these methods also utilize unique and prominent features allowing the medical images to be registered. Depending on the scale of the object, different reconstruction algorithms are applied to the data collected from the imaging process.

Feng and Ip [39] performed surface reconstruction on brain magnetic resonance imaging (MRI) data. Each region of the brain is approximated using a low-resolution geometric model. The model is then deformed to statistically and physically fit the high-resolution imaging data. The statistical and physical constraints allow the surface edges to be more robust to noise near the region boundaries. The low-resolution step allows the model to represent the large-scale structures from a small sample size of training subjects. Building a model from small training sets is particularly important for multi-dimensional reconstruction as high dimensional subject data is rare and the high dimensionality requires large amounts of memory.

When performing gross medical imaging, a series of 2D images might be stacked together if the object motion and the imaging system motion is negligible compared to the overall dimensions of the object. For example, when performing ultra-sound to image large organs, the vibrations of the ultra-sound probe and the small motion of muscles surrounding the organs are insignificant due to the relative scale of the object being imaged [22]. In addition, stacking can be acceptable if a stationary object reference is visible in each frame. When performing a brain MRI, the stationary bone structure of the skull can be used to translate the 2D scans for the registration process [53]. These methods are suitable for large scale stationary objects, and can be used to supplement small scale objects with ambitious features as long as other methods are also utilized.

On the smaller scale, electron microscopy is used to image cells. In these cases, the vibrations and motions of the cells are significant. However, like in gross medical imaging, electron microscopy can use reference points that are present in multiple 2D images [65].

Single-particle analysis [46, 52] attempts to identify macromolecules in each view and attempts to determine the orientation of each macromolecule in the particle.

Ahmad et. al [2] present research on volumetric reconstruction of bone density from four bone density scans, in the form of dual-energy x-ray absorptiometry (DXA). A 3D reconstruction was iteratively generated such that the bone density difference of 2D projections of the reconstruction onto the DXA imaging plane with the DXA data was minimal. Ahmad stated that bone density measurements sampled from their DXA reconstruction could replace qualitative computed tomograms (QCT), thereby reducing the radiation dose by a factor of 100. Like Ahmad, the reconstruction method proposed in this document also attempts to reduce the reconstruction error projected into the corneal tomogram domain. Existing corneal literature did not attempt to minimize reconstruction error.

Since retinal OCT imaging is similar to corneal imaging, it is useful to briefly examine Wojtkowski's 3D OCT retinal imaging system [140]. Wojtkowski created a volume rendering of a 3D region of size 6 $mm \times$ 6 $mm \times$ 1 mm and performed retinal boundary layer measurements on the result. There is great appreciation that the authors correctly claimed to create a volume rendering instead of a reconstruction since the retinal measurements were mapped directly into a volumetric rendering. Although the retinal thickness measurements looked promising, unfortunately the methods do not transfer to corneal layer thickness measurements, which will be shown in Subsection 2.2.5.

Findings

The non-corneal reconstructions utilize models that incorporate structure or pixel intensity similarity and all attempt to minimize reconstruction error. Given a corneal model, the reconstruction process should be able to produce a corneal reconstruction that accounts for anatomy and physiology. Further a corneal model should relax the dense sampling requirement that is required in stacking to populate voxel values.

2.2.5   Corneal layer boundary localization

Subsection 2.2.1 explains the importance of corneal anatomy in the vision process in addition to introducing the corneal structures and corneal layers that, if diseased or disordered, can prevent light from projecting correctly onto the retina. In particular, the shape, measured through corneal layer boundary locations, is a metric of disease and disorder. Since the goal of this document is to create a 3D reconstruction of the cornea, the corneal model

requires understanding of the layer boundary structure. This subsection provides background on methods that can locate these layer boundary layers within UHROCT samples. A statistical shape model can then be constructed using labelled boundaries obtained from UHROCT as data samples. This subsection also motivates the need for automated corneal layer boundary localization, which is proposed in Chapter 5.

Object boundary extraction in digital images has been a long-standing problem in the computer vision community. A deformable model, first formulated by Kass et al. [59], is an elegant approach and it and its variants are widely used for object boundary extraction. Deformable model-based image segmentation approaches are broadly divided into three categories: parametric [21, 59, 143], non-parametric [17], and user-interactive [85, 90], where the parametric and non-parametric methods require no user interaction, and the interactive methods prompt the user during the segmentation process. These three approaches model the object boundary using an energy minimizing spline and express the total energy as a sum of internal energy (prior) and external energy (measurement). Further, these approaches follow a specified optimization technique, such as gradient descent, to obtain the global minimum of the total energy, based on a curve evolution technique.

There is a fundamental problem with parametric and non-parametric active contours that limit their direct use for corneal layer extraction. Active contours are curves designed to surround lines and shapes that may be present in the image [3, 21, 59, 76]. The active contour converges when the sum of internal (prior) and external (measurement) forces are minimized, such that the internal forces prefer contour smoothness (or some other prior shape), and the external forces prefer a fit to the given image, normally related to the image gradient. Both parametric and non-parametric active contours are sensitive to local minima, and the cornea images obtained in this study are highly contaminated with speckle noise, which creates false local minima, potentially misleading the active contour towards false layers. Therefore, existing unsupervised parametric and non-parametric active contours are not suitable for corneal layer extraction. However existing user interactive active contours, such as intelligent scissors [85, 90], can be used for interactive segmentation.

Mortenson et al. [90] first introduced intelligent scissors (IS), a semi-automatic segmentation algorithm, a process that takes advantage of user knowledge when finding a boundary. For this method, the user selects seed points on the boundary and, as the mouse moves along the boundary, the optimal boundary path between the starting point and the current point is obtained using a dynamic programming method and hard constraints. There are two principal advantages to this approach. First, rapid image segmentation can be accomplished in real-time as the operator can trace the layer boundaries. Second, the boundary accuracy of interactive segmentation is generally higher than automatic methods since user knowledge is utilized throughout the process [90].
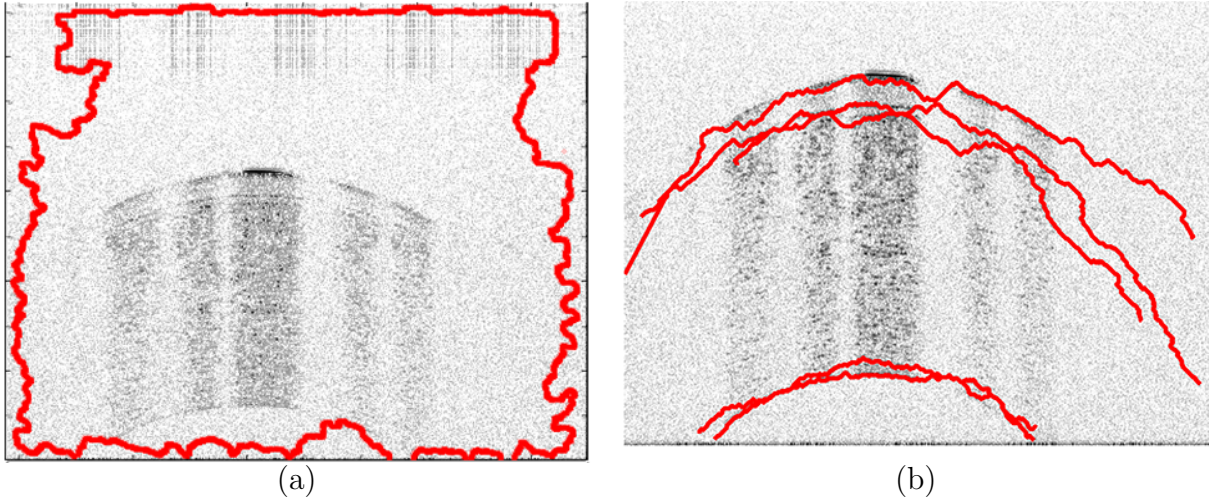
Figure 2.10: UHROCT segmentation results for (left) geometric active contour and (right) intelligent scissors. Neither method produces accurate segmentation results.

However, there are drawbacks to the IS methods for the corneal problem considered here. First, the boundary definition for IS methods relies on image gradients, making the algorithm sensitive to contrast non-uniformities found in low signal-to-noise (SNR) UHROCT images of the human cornea. Fig. 2.10 illustrates the performance of the algorithm despite having the user generate 20 to 30 points for each layer. Second, existing IS methods require the clinician to perform relatively accurate manual tracing along the region boundary, which is time consuming and laborious, particularly for complex regions of interest and large sets of three-dimensional corneal images, which may number in the hundreds or thousands.

Mishra et al. [85] proposed enhanced intelligent scissors (EIS) which uses a phase-based representation of the image as the external local cost, instead of the image gradient, and employs a Viterbi search to avoid the hard constraints of IS. However, as shown in Fig. 2.11, EIS produces non-smooth boundaries for the outer corneal layers and cannot distinguish Bowman's membrane and Descemet's endothelium complex due to insufficient contrast. EIS still requires the user to select many points on each layer boundary requiring about five to ten minutes for our trained technicians to *semi-automatically* segment a single 2D UHROCT image. While EIS is an improvement over completely manual segmentation, the resulting segmentation is sensitive to image artifacts and noise near the layer boundaries. Active contours were also investigated to automate the process, however preliminary testing on 20 real world tomograms showed that the active contour produced completely unusable

41

Figure 2.11: Enhanced intelligent scissors (EIS) [85] provides inaccurate localization of the two outer corneal boundaries and a third boundary that unpredictably weaves between the epithelium-Bowman boundary and the Bowman-stroma boundary. Further, Descemet's endothelium complex cannot be segmented.

results and was unable to segment the cornea from the background. Instead, to precisely locate the layers of the cornea, Chapter 5 proposes a novel model-based segmentation method that is automatic.

While many well developed retinal OCT imaging techniques exist to identify layer boundaries of the retina, corneal imaging provides different challenges. Garvin proposes the use of a general graph-based approached that attempts to reconstruct the retinal images into a 3D model and isolate the surfaces that correspond to the retinal layers [48]. In addition, Mishra developed a method using image gradient information and a kernel function to successfully compensate for the speckle noise, present in OCT images, and efficiently segments retinal layers [84]. The major difference between retinal and corneal segmentation is due to composition of the layers. As shown in Fig. 2.12, unlike the cornea, the retinal layers have different mean intensities for each layer. The retinal methods are good at finding the edge between these layers. In contrast, corneal layers have a similar mean intensity, but are separated by low-contrast, discontinuous, thin layer boundaries instead. As a result, retinal methods are able to find the high contrast outer layers, but could not locate the inner layers.

For corneal boundary localization, LaRocca et al. [67] use dynamic programming to isolate three of the five corneal layer boundaries on OCT images with an axial resolution of 3.4 μm, but do not differentiate the endothelium-Bowman's membrane boundary from the Bowman's membrane-stroma boundary.

Although UHROCT can resolve the five corneal layers, the endothelium has a layer thickness of only 5 μm and cannot be resolved with sufficient contrast from a single to-

Figure 2.12: Comparison of retinal (left) and corneal (right) UHROCTs. Unlike corneal layers, each retinal layer has a visibly distinct intensity compared to adjacent layers. In contrast, corneal layers contain a visible thin, dark boundary between each layer.

mogram on a 3 μm system. Instead the Descemet-endothelium complex, combining the Descemet's membrane and the endothelium layer, can be resolved. An automated segmentation method is presented in Chapter 5 to allow each layer boundary to be automatically identified. This method has been published as the first method to automatically identify the inner corneal boundaries and is shown to be capable of segmenting five corneal layer boundaries [35]. This method is useful to accurately measure the thickness of each layer so that a statistical 3D corneal model can be generated.

Chapter 3

Problem Formulation

This chapter defines the motivation, objectives, and scope required to better incorporate anatomy and physiology into general statistical reconstruction. Section 3.2 overviews all of the objectives that are introduced in this chapter and pursued in the following chapters. Subsection 3.2.1 describes these changes in relation to the general medical imaging applications. Subsection 3.2.2 applies these techniques to corneal imaging in hopes to create awareness to existing corneal reconstruction limitations and to provide a method to produce more meaningful reconstructions to the field.

To best understand this chapter, the reader should have familiarity with all of the high level components, notation, and concepts in general reconstruction, Section 2.1. The reader should also be familiar with existing corneal reconstruction methods, Subsection 2.2.4.

3.1   Existing limitations

From the discussion in Chapter 1 and Section 2.2, corneal reconstruction limitations were identified. Existing corneal reconstruction methods simply create a 3D field and populate the values directly from densely sampled measurements. These methods require hundreds of corneal tomograms that are then translated and stacked. Since a corneal model is not utilized the UHROCT measurement uncertainty due to imaging noise or eye-motion is not being quantified as part of the reconstruction process. While other medical imaging applications contain unique reference points that can be used for registration, UHROCT corneal tomograms have similar structural appearance from many different perspectives making it difficult to determine the location of the imaging plane.

While general reconstruction theory works very well for decoupled and linearized systems, it is challenging to formulate the general Bayesian estimator for the cornea. The FDOCT LCI imaging forward model, (2.49), to relate scattering potential measurements to cellular structure is highly non-linear [89]. Further the scattering potential measurements are spatially coupled to the distribution of cells within the cornea. If the measurement and corneal model state parameters can be decoupled and made into a linear system then reconstruction can become straightforward. These challenges probably result in the simplified one-to-one measurement-to-volume mapping common in existing state-of-the-art corneal reconstruction methods.

3.2   Objectives

The objectives introduced in Section 1.3 are listed below for reference. The details for Objective 1 are presented in Subsection 3.2.1 and the corneal specific Objectives 2 through 7 are detailed in Subsection 3.2.2.

1. Extending general reconstruction theory

   (a) Decouple subject model states
   (b) Estimate forward model parameters

2. Create 2D corneal structure model

3. Create 3D corneal model

   (a) Create corneal structure model
   (b) Create corneal texture model
   (c) Create corneal motion model

4. Generate synthetic cornea from model

5. Reconstruct cornea from tomograms without tomogram motion correction

6. Reconstruct cornea from tomograms with tomogram motion correction from pose estimates

7. Extract corneal layer thickness measurements

3.2.1 Extending general reconstruction theory

Section 2.1 presents a general theory used to reconstruct an arbitrary object from a given set of measurements. Due to the general nature of theory, there are additional challenges that are not addressed because they are specific to an application. For corneal imaging applications, imaging constraints and coupled subject models are common; the measurements are subject to the anatomy, physiology, and presence of pathologies. For corneal imaging, it is important to introduce additional notation and strategies into the existing statistical reconstruction theory.

In addition, physiological effects, such as motion, affect the imaging system. While statistical reconstruction works well when a clear correspondence between measurements and their respective properties in the subject model exists, any motion or change in subject behaviour make the correspondence uncertain. Therefore, in the presence of physiological effects, it is important to complement reconstruction with a physiological model and corresponding parameter estimation; when estimating subject model states the forward model parameters should also be estimated.

Objective 1, expands upon Section 2.1 by introducing corneal imaging specific components into the general theory. This objective can be broken down further into two sub-objectives: subject model decoupling and tomogram pose estimation.

Subject model decoupling

The first sub-objective is a strategy that can decouple the subject model using knowledge of anatomy and physiology. While the same strategy can be extended to pathology, the scope of the work is limited to healthy subjects. The key contribution lies in the fact that the subject model states, Z, can be rewritten as a, generally non-linear, decoupling function, ℓ, and a set of decoupled states, $\{Z_1, Z_2, \ldots\}$,

$$Z = \ell\left(Z_1, Z_2, \ldots\right). \tag{3.1}$$

Then a preprocessing function can be introduced and the forward model can be redefined such that,

$$\wp_1\left(\mathbf{m}\right) = f_1\left(Z_1\right) \tag{3.2}$$

$$\wp_2\left(\mathbf{m}\right) = f_2\left(Z_2\right) \tag{3.3}$$

$$\vdots, \tag{3.4}$$

with the goal to make Z straightforward to estimate from $\wp\,(\mathbf{m})$. The introduction of ℓ and $\wp$ make an otherwise convoluted reconstruction feasible. While Section 4.2 still provides a general description of how this can be applied to arbitrary medical imaging applications, Chapter 6 provides a specific example of how decoupling is applied specifically to UHROCT corneal imaging.

Corneal image pose estimation

The second sub-objective addresses the effects of physiology on measurement location uncertainty. Instead of estimating the measurement location corresponding to every pixel within a medical image, the scope of work is limited to estimating the location and orientation of each medical image. Many medical images are designed to rapidly sample a region of the subject such that any motion introduced by physiology is negligible, through a combination of high speed sampling and reducing physiological motion by influencing the subject. For this scope, it is assumed that physiological motion within a single image is negligible, but that effects of physiological motion are apparent between the acquisition of sequential images. Section 6.3 verifies these assumptions for the corneal reconstruction application. Section 4.1 introduces a set of parameters, $\boldsymbol{\Theta}$ into the forward model,

$$\mathbf{m} = f\left(Z, \boldsymbol{\Theta}\right), \tag{3.5}$$

that define the pose and center of the imaging plane. Further, Section 4.3 presents a general optimization formulation that attempts to estimate $\boldsymbol{\Theta}$ for each medical image through reconstruction error minimization.

3.2.2   Corneal reconstruction

As shown in Subsection 2.2.1, the central corneal thickness is a useful statistic in determining corneal health and the effectiveness of treatment. However a better metric is to also include the thickness of each corneal layer instead of only measuring the thickness of the entire cornea. However locating each corneal layer within an UHROCT tomogram is time consuming since it requires an expert operator to manually segment and measure each corneal layer boundary. An automated corneal layer segmentation tool, Objective 2, would allow the corneal layer boundaries to be extracted and measured allowing practitioners and researchers to immediately obtain thickness measurements without having to wait for the manual segmentation expert to be available. Further, the automated segmentation tool can be applied to thousands of tomograms providing statistics for a corneal model. This tool is created in Chapter 5.

As shown in Subsection 2.2.4, existing medical reconstruction techniques can create 3D corneal volumes by correlating multiple 2D tomograms and stacking them to create a volume. While correlation is useful in retinal images containing unique features within a tomogram, UHROCT corneal tomograms have layer boundaries that have similar appearance in every tomogram. The lack of distinct features combined with eye motion make it difficult for the operator to determine precisely from where the tomogram is sampled. While the stacking approach will produce a nice looking 3D volume, it does not take into account any physiological eye-motion, it does not take advantage of known corneal anatomy during the reconstruction, it does not have a model from which to estimate scattering potential in unsampled regions of the cornea or to reduce measurement noise in densely sampled regions.

Corneal model decoupling

Chapter 6 and Chapter 7 attempt to address these limitations by constructing a corneal model, Objective 3, and by reconstructing a cornea through corneal model state estimation, Objective 5. Since UHROCT measures scattering potential, $\mathbf{m}$, at a specified measurement location, a convenient corneal model can simply contain a field of scattering potentials, Z, that correspond to $\mathbf{m}$ through a forward model reduced to a correspondence matrix, C, such that

$$\mathbf{m} = C\mathbf{z}. \tag{3.6}$$

The matrix C is determined by mapping the measurement from the imaging plane into a point on the cornea in 3D space, derived in Section 4.1. Existing techniques use pixel intensity correspondence to estimate C and, without a corneal model, guesses at values for Z that do not have a direct corresponding $\mathbf{m}$.

By decoupling Z into two sets of states, Z_1 and Z_2, a corneal anatomy model and a scattering potential model can be created and used to estimate Z based on the value of every single measurement, allowing regions of Z to be estimated without a corresponding $\mathbf{m}$. The states in Z_1 correspond to parameters in the corneal anatomy model used to describe the shape of each corneal layer, Section 6.1. The states in Z_2 correspond to scattering potential. With the structure removed, Z_2 can be modelled as a stationary Markov random field. This allows the neighbouring relationships within Z_2 to be modelled as a kernel determined through analysis of the scattering potential found in a tomogram, Section 6.2. The statistical model of Z_2 can be represented as a state covariance matrix, $\mathcal{P}_2$. The motion of the eye is also modelled in Section 6.3 but is used after the reconstruction process when attempting to estimate tomogram pose in Chapter 8.

Synthetic data

Using the corneal models, synthetic data can also be generated since it provides an unlimited source of ground truth for use in evaluating reconstruction and tomogram estimation. Objective 4 is satisfied throughout Chapter 6 and the resulting synthetic dataset is described in Section 7.1.

Reconstruction

During reconstruction, measurement preprocessing can determine the location of each corneal layer boundary using the automated method created in Chapter 5. Let $\wp_1$ be the preprocessing function that generates a vector of states that correspond to Z_1 and let $\wp_2$ be the preprocessing function that generates a vector of scatter potential that has been mapped into a random field such that the corneal structure has been removed,

$$\wp_1\left(\mathbf{m}\right) = C_1 Z_1 \tag{3.7}$$

$$\wp_2\left(\mathbf{m}\right) = C_2 Z_2. \tag{3.8}$$

Then Z_2 can be estimated using Bayesian linear least-squares estimation [42, 123] and an estimation of measurement noise, R,

$$\hat{Z}_2 = \left(C_2^T R^{-1} C_2 + P_2^{-1}\right) C_2^T R^{-1} \left(\wp_2\left(\mathbf{m}\right) - \mu_{Z_2}\right) + \mu_{Z_2}. \tag{3.9}$$

The value of Z_1 can be found using $\wp_2\left(\mathbf{m}\right)$ and Z_2 through surface parameter estimation described in Section 7.2.

In addition to a 3D volume of scattering potential data, Z, and is also obtained using existing corneal reconstruction techniques, the reconstruction also produces estimates for all of the corneal layer boundaries, Z_1. This allows the operator to easily obtain corneal boundary layer thickness measurements at any arbitrary cross section, satisfying Objective 7.

Pose estimation

The final object is to improve pose estimation using the corneal model, Objective 6. In Chapter 8, the translation and orientation of each corneal tomogram are estimated using the optimization formulation in Section 4.3 and previously described in Subsection 3.2.1. Since C is derived by transforming the location of each tomogram measurement into a

position on the cornea, C is dependent on the pose of the tomogram, and can be written explicitly as $C(\boldsymbol{\Theta})$,

$$\mathbf{m} = C\left(\boldsymbol{\Theta}\right)\mathbf{z}, \tag{3.10}$$

where $\boldsymbol{\Theta}$ contains all of the parameters to map the location of $\mathbf{m}$ to the location in Z. The values of $\boldsymbol{\Theta}$ are determined through an iterative optimization formulation that minimizes the reconstruction error,

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \left| C(\boldsymbol{\Theta})Z(\hat{\boldsymbol{\Theta}}) - \mathbf{m} \right|. \tag{3.11}$$

Notice that $\hat{Z}$ is also a function of $\boldsymbol{\Theta}$ because it implicitly contains $C(\boldsymbol{\Theta})$.

The optimization cost function becomes slightly more complicated when the decoupled states are introduced,

$$\hat{Z}(\boldsymbol{\Theta}) = \ell\left(\hat{Z}_1(\boldsymbol{\Theta}), \hat{Z}_2(\boldsymbol{\Theta})\right) \tag{3.12}$$

$$\hat{\boldsymbol{\Theta}} = \arg\min_{\boldsymbol{\Theta}} \left| C(\boldsymbol{\Theta})\ell\left(\hat{Z}_1(\boldsymbol{\Theta}), \hat{Z}_2(\boldsymbol{\Theta})\right) - \mathbf{m} \right|. \tag{3.13}$$

Again, the estimation of Z_1 and Z_2 are dependent on $\boldsymbol{\Theta}$ since the tomogram pose affects the reconstructed corneal surface, C_1 and, similarly, how the scattering potential is mapped, C_2.

Chapter 4

Extended Statistical Reconstruction

This chapter focuses on extending statistical reconstruction theory with additional notation and components specific to reconstruction involving a set of medical images, or tomograms T, that are obtained from a subject. The values of T are measurements of unknown subject states, Z. If Z contains coupled states, it can be difficult to estimate Z from T. For example, in corneal reconstruction, the scattering potential measurements are coupled with the position of the measurement within the cornea making it difficult to ascertain if a difference in scattering potential is due to a change in cellular density or a change in the amount of backscatter when imaging near a corneal layer boundary. While a human OCT operator can learn to decouple position from cellular density through tomogram observation, a corneal reconstruction algorithm requires a mathematical model. The most challenging aspect of corneal reconstruction is creating a model that decouples Z and relates T to Z linearly.

Once Z is decoupled, a forward model, f, can contain unknown parameters. For corneal reconstruction, these unknown parameters are the positions where measurements are sampled within the cornea. Since the subject moves during imaging, the desired measurement location and the actual measurement location differ, creating measurement location uncertainty. The pose, position and orientation, of each tomogram at the time of measurement needs to be estimated in addition to the states in Z, in order to reduce measurement location uncertainty. Section 4.3 describes a means to reconstruct the cornea while also estimating the pose of each tomogram. Chapter 8 tests this theory during corneal reconstruction.

Chapter organization

This chapter is divided into three sections. Section 4.1 defines a linear relationship between the measurement locations in T to their locations in 3D space in the actual cornea. This section introduces notation required to express the location of measurement in T in 3D space given an arbitrary pose for T. Section 4.2 provides a brief overview of the decoupling strategy that is used throughout Chapter 6. The last section, Section 4.3, describes a method to estimate the pose of each tomogram using an optimization formulation.

4.1 Measurement mapping

This section describes how the locations of each measurement within a set of tomograms, T, are mapped to the corresponding sampling location in 3D space.

4.1.1 Lexicographical reordering

The tomograms are mapped into a vector so that multi-dimensional matrices can be manipulated using normal matrix operations. Using lexicographical reordering, (2.6), T_i can be mapped into $\mathbf{m}_i$,

$$\mathbf{m}_i = T_{i,:}, \tag{4.1}$$

which can be rewritten in terms of pixel k, the associated measurement, $m_{i,k}$, and the position of pixel k within T_i, $\mathbf{s}_{i,k}$,

$$k = s_{0,i,k} + s_{1,i,k}\eta_{s,0,i}, \tag{4.2}$$

$$m_{i,k} = T_i(\mathbf{s}_{i,k}). \tag{4.3}$$

Using the relationship between $\mathbf{m}$ and $\mathbf{m}_i$ from (2.1), $\mathbf{m}$ can contain all of the measurements from T,

$$\mathbf{m} = \begin{bmatrix} \mathbf{m}_0^T & \mathbf{m}_1^T & \mathbf{m}_2^T & \cdots & \mathbf{m}_i^T & \cdots \end{bmatrix}^T. \tag{2.1 on page 9}$$

4.1.2 Planar notation

Tomogram i, T_i, contains a set of measurements, $\mathbf{m}_i$, that were sampled within an imaging plane Π_i. The plane Π_i is centered at the point $\mathbf{c}_i$ in $\Re^3$,

$$\mathbf{c}_i = \begin{bmatrix} c_{0,i} & c_{1,i} & c_{2,i} \end{bmatrix}^T, \tag{4.4}$$

and has a normal vector denoted as $\vec{\mathbf{n}}_i$. Since the image on the tomogram can be rotated arbitrarily within the plane, θ_i represents the amount that T_i is rotated about the axis $\vec{\mathbf{n}}_i$, as illustrated in Fig. 4.1. Incorporating $\gamma_{s,i}$, the scale of T_i defined in (2.34), all of the parameters required to represent the scale and pose, the position and orientation, of T_i are collected into the vector $\boldsymbol{\Theta}_i$, let

$$\boldsymbol{\Theta}_i \triangleq \begin{bmatrix} \vec{\mathbf{n}}_i \\ \mathbf{c}_i \\ \theta_i \\ \gamma_{s,i} \end{bmatrix}. \tag{4.5}$$

The parameters for all tomgrams in the set, T, are denoted using $\boldsymbol{\Theta}$,

$$\boldsymbol{\Theta} = \begin{bmatrix} \boldsymbol{\Theta}_1 \\ \boldsymbol{\Theta}_2 \\ \vdots \\ \boldsymbol{\Theta}_i \\ \vdots \end{bmatrix}. \tag{4.6}$$

4.1.3 Define $\Re^3$ notation

Let $Z \in \Re^3$ be spanned by the basis vectors, $\vec{\mathbf{x}}_0$, $\vec{\mathbf{x}}_1$, and $\vec{\mathbf{x}}_2$,

$$\vec{\mathbf{x}}_0 = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T \tag{4.7}$$

$$\vec{\mathbf{x}}_1 = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}^T \tag{4.8}$$

$$\vec{\mathbf{x}}_2 = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}^T. \tag{4.9}$$

The arbitrary point, $\mathbf{x}$ in $\Re^3$, is positioned x_0 along $\vec{\mathbf{x}}_0$, x_1 along $\vec{\mathbf{x}}_1$, and x_2 along $\vec{\mathbf{x}}_2$

$$\mathbf{x} = x_0 \cdot \vec{\mathbf{x}}_0 + x_1 \cdot \vec{\mathbf{x}}_1 + x_2 \cdot \vec{\mathbf{x}}_2. \tag{4.10}$$

Figure 4.1: The random field Z (blue) is sampled using the planar subspace of tomogram T_i (orange). Let $\mathbf{c}_i$ (red) be the center position of T_i, let $\overrightarrow{\mathbf{n}}_i$ (green) be the vector normal to the sampling plane, and let θ_i (black) be the rotation of the plane about $\overrightarrow{\mathbf{n}}_i$ through $\mathbf{c}_i$.

Let $\eta_{x,0,j}$, $\eta_{x,1,j}$, and $\eta_{x,2,j}$,

$$\eta_{x,j} = \begin{bmatrix} \eta_{x,0,j} & \eta_{x,1,j} & \eta_{x,2,j} \end{bmatrix}^T, \tag{4.11}$$

represent the number of elements in the $\vec{\mathbf{x}}_0$, $\vec{\mathbf{x}}_1$, and $\vec{\mathbf{x}}_2$ directions, respectively and let $\gamma_{x,j}$ represent the dimensions of Z. Let $\mathbf{x}_j$ represent the position of z_j in $\Re^3$ and let $x_{0,j}$, $x_{1,j}$, and $x_{2,j}$ represent the number of elements in the $\vec{\mathbf{x}}_0$, $\vec{\mathbf{x}}_1$, and $\vec{\mathbf{x}}_2$ directions. The center of field Z is located at $\mathbf{c}_j$ in $\Re^3$. See Fig. 4.1 for an illustration of these parameters.

4.1.4 Coordinate mapping

A mapping function is needed to convert from $\mathbf{s}$ in $\Re^2$ to $\mathbf{x}$ in $\Re^3$ to establish a spatial correspondence, C, in Subsection 4.1.5 between $\mathbf{m}$ and $\mathbf{z}$. Fig. 2.6 illustrates the arbitrary point $T_i(\mathbf{s})$ in the local image space of T_i. The horizontal and vertical directions in T_i corresponds to the directional vectors $\vec{\mathbf{s}}_{0,i}$ in $\Re^3$ and $\vec{\mathbf{s}}_{1,i}$ in $\Re^3$, respectively. The vectors $\vec{\mathbf{s}}_{0,i}$ and $\vec{\mathbf{s}}_{1,i}$ can be found in terms of $\boldsymbol{\Theta}_i$, by finding two mutually orthonormal vectors orthogonal to $\vec{\mathbf{n}}_i$ using a series of cross-products. When performing higher dimensional reconstructions, the Gram-Schmidt process [113] can be utilized to find these mutually orthonormal vectors more efficiently.

The following derives a change of basis that allows the point $\psi \in \Re^2$ to be transformed into $\mathbf{x} \in \Re^3$. The details of the derivation are provided for completeness and the results are necessary to determine a correspondence between $\mathbf{m}$ and Z in Chapter 7. Equations (4.18), (4.19), and (4.20) can be accepted and the derivation can be skipped on the first read through without loss of understanding in other chapters.

The first basis vector, $\mathbf{u}_{0,i}$, is equal to $\vec{\mathbf{n}}_i$,

$$\mathbf{u}_{0,i} = \vec{\mathbf{n}}_i. \tag{4.12}$$

The second basis direction, $\mathbf{u}_{1,i}$, is found by calculating the cross product of $\mathbf{u}_{0,i}$ with $\vec{\mathbf{x}}_0$, $\vec{\mathbf{x}}_1$, and $\vec{\mathbf{x}}_2$,

$$\mathbf{u}_{1,i,x_{0,i}} = \mathbf{u}_{0,i} \times \vec{\mathbf{x}}_0, \qquad \mathbf{u}_{1,i,x_{1,i}} = \mathbf{u}_{0,i} \times \vec{\mathbf{x}}_1, \qquad \text{and } \mathbf{u}_{1,i,x_{2,i}} = \mathbf{u}_{0,i} \times \vec{\mathbf{x}}_2. \tag{4.13}$$

Then, the most numerically stable vector, defined as the vector with the greatest L^2-norm, either $\mathbf{u}_{1,i,x_{0,i}}$, $\mathbf{u}_{1,i,x_{1,i}}$, or $\mathbf{u}_{1,i,x_{2,i}}$, is selected as the second basis vector. The third basis direction, $\mathbf{u}_{2,i}$, can be calculated using the cross-product,

$$\mathbf{u}_{2,i} = \mathbf{u}_{0,i} \times \mathbf{u}_{1,i}. \tag{4.14}$$

The final step is to define the basis vectors, $\mathbf{e}_{0,i}$, $\mathbf{e}_{1,i}$, and $\mathbf{e}_{2,i}$ by normalizing $\mathbf{u}_{0,i}$, $\mathbf{u}_{1,i}$, and $\mathbf{u}_{2,i}$,

$$\mathbf{e}_{0,i} = \frac{\mathbf{u}_{0,i}}{\|\mathbf{u}_{0,i}\|} \qquad e_{1,i} = \frac{\mathbf{u}_{1,i}}{\|\mathbf{u}_{1,i}\|}, \qquad \text{and } \mathbf{e}_{2,i} = \frac{\mathbf{u}_{2,i}}{\|\mathbf{u}_{2,i}\|}. \qquad (4.15)$$

Then, since the imaging plane parameter θ_i allows the entire plane to rotate about $\overrightarrow{\mathbf{n}}_i$, $\mathbf{e}_{1,i}$ and $\mathbf{e}_{2,i}$ are rotated by θ_i radians about $\mathbf{e}_{0,i}$ allowing the horizontal and vertical directions in T_i to be mapped to vectors $\overrightarrow{\mathbf{s}}_{0,i}$ and $\overrightarrow{\mathbf{s}}_{1,i}$,

$$\overrightarrow{\mathbf{s}}_{0,i} = \mathbf{e}_{1,i} \cos \theta_i + \mathbf{e}_{2,i} \sin \theta_i \qquad (4.16)$$

and

$$\overrightarrow{\mathbf{s}}_{1,i} = \mathbf{e}_{2,i} \cos \theta_i - \mathbf{e}_{1,i} \sin \theta_i. \qquad (4.17)$$

Given that T_i consists of $\eta_{s,0,i}$ by $\eta_{s,1,i}$ pixels and has a width and height of $\gamma_{s,0,i}$ by $\gamma_{s,1,i}$ μm, the mapping from $\mathbf{s}_i$ to $\mathbf{x}_i$ is represented as

$$\mathbf{x}_i = \left(s_{0,i} - \frac{\eta_{s,0,i}}{2} \right) \frac{\gamma_{s,0,i}}{\eta_{s,0,i}} \overrightarrow{\mathbf{s}}_{0,i} + \left(s_{1,i} - \frac{\eta_{s,1,i}}{2} \right) \frac{\gamma_{s,1,i}}{\eta_{s,1,i}} \overrightarrow{\mathbf{s}}_{1,i} + \mathbf{c}_i. \qquad (4.18)$$

Fig. 4.2 illustrates the transformation. Finally, using $\mathbf{x}_i$, the positional correspondence between pixel k on tomogram i and the sampling position within the volume Z can be defined as

$$\mathbf{x}_{i,k} = \left(s_{0,i,k} - \frac{\eta_{s,0,i}}{2} \right) \frac{\gamma_{s,0,i}}{\eta_{s,0,i}} \overrightarrow{\mathbf{s}}_{0,i} + \left(s_{1,i,k} - \frac{\eta_{s,1,i}}{2} \right) \frac{\gamma_{s,1,i}}{\eta_{s,1,i}} \overrightarrow{\mathbf{s}}_{1,i} + \mathbf{c}_i. \qquad (4.19)$$

The forward model for the entire imaging plane, can be rewritten for pixel k on tomogram i,

$$T_i(\mathbf{s}_{i,k}) = f\left(Z, \mathbf{x}_{i,k} \right) \qquad (4.20)$$

These two equations are essential for reconstruction since (4.19) is required to map the positions of measurements from T to their corresponding positions in Z, in the following section, and (4.20) defines the forward model used in Subsection 6.3.2.

Figure 4.2: This illustration summarizes the mapping from (left) $T_i(\mathbf{s}_i)$, a point in the image space, to (right) $\mathbf{x}_i$, the point in Z space.

4.1.5 Correspondence matrix

Using the derivation from Subsection 4.1.4, a correspondence matrix, C, that maps the location of measurements within a tomogram and the corresponding location of measurements in the cornea can be expressed. Although the mapping equations can appear algorithmic, rest assured that C is very useful in Chapter 7 and Chapter 8 when $\mathbf{m}$ is mapped into Z during reconstruction.

Then $\mathbf{x}_j$ can be determined in terms of $x_{0,j}$, $x_{1,j}$, and $x_{2,j}$,

$$j = x_{0,j} + x_{1,j}\eta_{x,0,j} + x_{2,j}\eta_{x,1,j}\eta_{x,0,j}, \tag{4.21}$$

$$\mathbf{x}_j = \begin{bmatrix} \frac{\gamma_{x,0,j}}{\eta_{x,0,j}}\left(x_{0,j} - \frac{\eta_{x,0,j}}{2}\right) + c_{0,j} \\ \frac{\gamma_{x,1,j}}{\eta_{x,1,j}}\left(x_{1,j} - \frac{\eta_{x,1,j}}{2}\right) + c_{1,j} \\ \frac{\gamma_{x,2,j}}{\eta_{x,2,j}}\left(x_{2,j} - \frac{\eta_{x,2,j}}{2}\right) + c_{2,j} \end{bmatrix}. \tag{4.22}$$

If the tomogram point $\mathbf{x}_{i,k}$ and the point in z_j, $\mathbf{x}_j$ correspond, matrix $C^i_{k,j}$ will have a value of 1, otherwise it will have a value of 0,

$$C^i_{k,j} = \begin{cases} 1 & \text{if all}\big(|\,\mathbf{x}_{i,k} - \mathbf{x}_j\,| < \gamma_{x,j}./\eta_{x,j}\big) \\ 0 & \text{else} \end{cases}. \tag{4.23}$$

The function all$(\dots)$ is true if all elements of $(\dots)$ are true and the function is false otherwise allowing $C^i_{k,j}$ to be 1 if the tomogram element position is contained within the z_j position. The operator $./$ is elementwise division. Finally, C can be determined by stacking all C^i matrices,

$$C = \begin{bmatrix} C^0 \\ C^1 \\ \vdots \\ C^i \\ \vdots \end{bmatrix}. \tag{4.24}$$

During the construction of C, values from (4.19) and (4.22) are substituted into (4.23).

4.2 Decoupling

The corneal model in Chapter 6 is designed to decouple the states Z so that reconstruction can estimate Z using a Bayesian linear least-squares estimator, which is preferable to the general Bayesian estimator. The decoupling strategy is outlined here and demonstrated in more detail in Chapter 6.

Given a forward measurement model, f, and a collection of measurements $\mathbf{m}$, a parametric subject model can be developed. The parameters of the model populate a random field, Z, such that $\mathbf{m} = f(Z)$. However structural and behavioural characteristics can decompose Z into a set of independent states, $\{Z_1, Z_2, \ldots\}$, related by the function ℓ,

$$Z = \ell(Z_1, Z_2, \ldots) \tag{4.25}$$

to ease estimation of Z from measurements $\mathbf{m}$ and a forward model f.

The key is to decouple Z into sub-models based on existing subject domain knowledge. For example, if general knowledge of corneal anatomy is known, then that knowledge can be parameterized and separated from other states within Z. By introducing prior knowledge into sub-models, the data become more constrained, creating fewer degrees of freedom that need to be estimated, and ensuring that the data behave according to physical or behavioural restrictions.

For example, structural constraints applied to layers within a human cornea can be modelled as part of ℓ; ℓ can decouple structural parameters, Z_1, from the visual scatter potential characteristics, Z_2, removing dependency of scattering potential characteristics from relative position to corneal structure. A Bayesian linear least-squares estimator can be applied to the resulting linear scattering potential model relating $\mathbf{m}$ and Z_2. Then a specifically designed non-linear structural estimator can be applied to a carefully designed structural model with parameters to determine Z_1 from $\mathbf{m}$. The reconstruction component estimates values of Z given f and $\mathbf{m}$, allowing unmeasured elements of Z to be estimated and reducing measurement noise through a prior subject model of Z.

4.3 Pose estimation

While imaging the cornea using UHROCT, a set of tomograms, T are obtained. While it is desired to sample tomogram i at imaging plane Π_i, eye-motion during imaging changes the actual location of Π_i. The purpose of tomogram pose estimation is calculate $\hat{\Theta}_i$, the estimate for the actual imaging plane parameters, Θ_i. By formulating an optimization

problem, $\hat{\boldsymbol{\Theta}}_i$ can be calculated for all tomograms with the goal of reducing tomogram pose error, $\tilde{\boldsymbol{\Theta}}_i$,

$$\tilde{\boldsymbol{\Theta}}_i = \boldsymbol{\Theta}_i - \hat{\boldsymbol{\Theta}}_i. \tag{4.26}$$

However, unless ground truth is available, the actual parameters in $\boldsymbol{\Theta}_i$ are not known, so the reconstruction measurement residual error is used instead,

$$\tilde{\mathbf{m}} = \mathbf{m} - \hat{\mathbf{m}}. \tag{2.29 on page 16}$$

The pose estimation problem should attempt to minimize the reconstruction error over the entire field, otherwise a single tomogram might improve at a time to the detriment of fitting all of the data to the field. For example, if only the reconstruction error specific to tomogram i is minimized, $\tilde{\mathbf{m}}_i$ can be driven to zero by deforming the corneal structure and scattering potential field causing all other tomograms to have very high reconstruction error. Consequently, the goal is to minimize the sum of squared error,

$$\min \sum (\tilde{\mathbf{m}})^2, \tag{4.27}$$

which leads to

$$\hat{\boldsymbol{\Theta}} = \arg_{\boldsymbol{\Theta}'} \min \sum_{\forall i, \forall k} \left(m_{i,k} - f\left(Z, \mathbf{x}_{i,k}(\boldsymbol{\Theta}'_{\mathbf{i}})\right) \right)^2, \tag{4.28}$$

by substitution. It is convenient to use the mean squared error metric since the values can be better compared across different subjects when the number of measurements differ. However minimizing either the sum of squared error or the mean squared error makes no difference as long as the number of measurements remain constant.

The mean squared error metric was selected over the mean absolute error metric for two reasons. First, the mean squared error criterion penalizes tomograms that are farther from their correct locations more heavily than the mean absolute error criterion. Pose estimation allows these misaligned tomograms, that would otherwise be outliers, to move so that a better fitting corneal surface can be constructed. The second reason is that during optimization, the derivative of the absolute value function has a discontinuity that requires additional complexity when working with derivatives of the minimization function.

In general, the approach to solving this optimization problem depends on f. A nonlinear forward model, coupled states, and a discrete coordinate system transformation dependent on $\boldsymbol{\Theta}_i$, can be difficult to solve. In the worst case, a very unlikely case, the

values of f can have large variability producing a field with very little correspondence to Z. Instead of attempting to solve (4.28) for the general case, a specific f leads to a much more meaningful and succinct discussion.

Assuming Chapter 6 can decouple Z into Z_1 and Z_2 such that,

$$\wp_1 (\mathbf{m}) = C_1 Z_1 \qquad\qquad (3.7 \text{ on page } 49)$$

$$\wp_2 (\mathbf{m}) = C_2 Z_2, \qquad\qquad (3.8 \text{ on page } 49)$$

with a slight modification to introduce $\boldsymbol{\Theta}$,

$$\wp_1 (\mathbf{m}) = C_1(\boldsymbol{\Theta})Z_1(\boldsymbol{\Theta}) \qquad\qquad (4.29)$$

$$\wp_2 (\mathbf{m}) = C_2(\boldsymbol{\Theta})Z_2(\boldsymbol{\Theta}), \qquad\qquad (4.30)$$

equation (4.28) can be rewritten as a weighted combination, using ϖ, of residual errors,

$$\hat{\boldsymbol{\Theta}} = \arg_{\boldsymbol{\Theta}'} \min \varpi \big(\wp_1 (\mathbf{m}) - C_1(\boldsymbol{\Theta}')\mathbf{z}_1(\boldsymbol{\Theta}')\big)^2 + (1 - \varpi)\big(\wp_2 (\mathbf{m}) - C_2(\boldsymbol{\Theta}')\mathbf{z}_2(\boldsymbol{\Theta}')\big)^2. \quad (4.31)$$

Although $C_1(\boldsymbol{\Theta}')$ and $C_2(\boldsymbol{\Theta}')$ are not linear in terms of $\boldsymbol{\Theta}'$, they are sparse matrices containing ones and zeros. Unfortunately, despite the simple appearance of C_1 and C_2, the construction involves coordinate transformations using many of the derivations from Section 4.1 required to express (4.24).

While it is possible to formulate $C(\boldsymbol{\Theta})\mathbf{z}(\boldsymbol{\Theta})$ as a continuous field relating each measurement to every single element in $\mathbf{z}$, eliminating the stepwise construction of C, the sparsity of C is lost and the resulting relationship, while differentiable, is still difficult to solve directly. For example, $m_{i,k}$ can be related to $\mathbf{z}$ by weighting every single value of $\mathbf{z}$ as a decaying exponential function of the distance $\mathbf{x}_{i,k}$ to $\mathbf{x}_j$ for all i, j, and k,

$$\hat{m_{i,k}} = \sum_j z_j(\boldsymbol{\Theta}') \cdot \exp\Big(-\alpha\big(\mathbf{x}_j - \mathbf{x}_{i,k}\big)^2\Big). \qquad\qquad (4.32)$$

Equation (4.32) becomes much more complicated after (4.19) is substituted.

Instead, an iterative optimization approach is preferred using an initial estimate of $\boldsymbol{\Theta}$, and the Jacobian [80], J, of $\hat{\mathbf{m}}$ for $\boldsymbol{\Theta}$ and its corresponding reconstruction, $\hat{Z}(\boldsymbol{\Theta})$,

$$J(\hat{\mathbf{m}}) = \begin{bmatrix} \frac{\delta \hat{\mathbf{m}}_{1,1}}{\delta \boldsymbol{\Theta}_1} & \frac{\delta \hat{\mathbf{m}}_{1,1}}{\delta \boldsymbol{\Theta}_2} & \cdots \\ \frac{\delta \hat{\mathbf{m}}_{1,2}}{\delta \boldsymbol{\Theta}_1} & \frac{\delta \hat{\mathbf{m}}_{1,2}}{\delta \boldsymbol{\Theta}_2} & \cdots \\ \vdots & \vdots & \ddots \\ \frac{\delta \hat{\mathbf{m}}_{i,1}}{\delta \boldsymbol{\Theta}_1} & \frac{\delta \hat{\mathbf{m}}_{i,1}}{\delta \boldsymbol{\Theta}_2} & \cdots \\ \frac{\delta \hat{\mathbf{m}}_{i,2}}{\delta \boldsymbol{\Theta}_1} & \frac{\delta \hat{\mathbf{m}}_{i,2}}{\delta \boldsymbol{\Theta}_2} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}, \qquad\qquad (4.33)$$

where $\frac{\delta \hat{m}_{i,k}}{\delta \Theta_{...}}$ can be estimated using a central difference numerical differentiation approximation,

$$\frac{\delta \hat{m}_{i,k}}{\delta \Theta_{...}} \approx \frac{C(\Theta + \Delta\Theta)\hat{\mathbf{z}}(\Theta) - C(\Theta - \Delta\Theta)\hat{\mathbf{z}}(\Theta)}{2\,\|\Delta\Theta\|}. \tag{4.34}$$

Note that central difference approximation [13] is used because the function C is evaluated at $C(\Theta + \Delta\Theta)$ and $C(\Theta - \Delta\Theta)$ to calculate the Hessian [80]. If a Hessian approximation is determined using quasi-Newton methods [49], then single-point numerical differentiation [13] in the Jacobian can be used instead.

Using the Jacobian and Hessian, or an approximation of the Hessian, existing trust-region methods [24, 88, 127] can be used to estimate the tomogram pose.

Chapter 5

Corneal Layer Boundary Localization

Since the corneal model requires precise measurements from which to form a statistical model, the design of a 2D corneal tomogram model and an automated layer boundary localization tool allows layer measurements to be quickly generated from UHROCT data. The results from the localization tool can be used to generate synthetic test images and to determine reasonable parameter estimates for a 3D model. An expert technician can determine the location of each boundary by observing faint layer boundaries due to the difference in refractive index for each layer, which is visible in an UHROCT tomogram, Fig. 2.5, as a subtle change in scattering potential intensity. Using this approach, an automated tool should also be able to segment the layers by determining the most likely position of each layer through the relative difference in scattering potential at their interfaces. A structural model can then be created using measured layer boundary thickness and the refractive index at their interfaces, with the assumption that the inner-layer cellular structure can be represented as a texture in Section 6.2.

Any curve, $\Omega(s)$ where $s \in \Re^1$, can be used to represent a layer boundary. The true layer boundaries are represented by the curve $\Omega_\alpha(s)$, where $\alpha \in [0,1]$. The index variable ι is introduced so that each true boundary can be represented as $\Omega_{\alpha_\iota}(s)$, where $\iota \in \{1,2,3,4,5\}$ and $\alpha_\iota \leq \alpha_{\iota+1}$, corresponding to the epithelium (α_1), Bowman's membrane (α_2), stroma (α_3), Descemet's-Endothelium complex (α_4), and the endothelium (α_5).

Since the inner layers are sandwiched between the extreme (outer) layers, the 2D tomogram model asserts that the shape of each inner layer can be expressed as a continuous transformation between the shape of the epithelium and endothelium boundaries, such that the true (unknown) curve $\Omega_\alpha(s)$ is approximated as $\hat{\Omega}_\alpha(s)$, where

$$\hat{\Omega}_\alpha(s) = \alpha \hat{\Omega}_1(s) + (1-\alpha)\hat{\Omega}_0(s). \tag{5.1}$$

and boundary layers $\hat{\Omega}_0(s)$ and $\hat{\Omega}_1(s)$ are defined by arbitrary parameterized curves, $\mathcal{C}(s, \Theta)$,

$$\hat{\Omega}_0(s) = \mathcal{C}(s, \Theta_0) \tag{5.2}$$

$$\hat{\Omega}_1(s) = \mathcal{C}(s, \Theta_1). \tag{5.3}$$

Each layer boundary is then approximated as a linear combination of the epithelium and endothelium boundaries. The curves $\mathcal{C}(s, \Theta_l)$ parameterize $\Omega_l(s)$, using some parameter set Θ_l, for $l \in \{0, 1\}$. Since the outermost corneal boundaries have high contrast and can be consistently and reliably segmented, $\mathcal{C}(s, \Theta_l)$ can be determined by locating the outermost boundaries. The curve $\mathcal{C}(s, \Theta_l)$ can be parameterized using a low-dimensional corneal model; ellipses, quadratics, and cubic polynomials were tested, however these curves were not able to capture all of the manually segmented corneal layer curvature throughout the entire OCT tomogram. For the given datasets, a fourth-order polynomial, $\mathcal{P}(s, \Theta_l)$, was found to be a suitable model for a healthy cornea:

$$\Theta_l = \begin{bmatrix} \theta_{l,0,0} & \theta_{l,0,1} & \theta_{l,0,2} & \theta_{l,0,3} & \theta_{l,0,4} \\ \theta_{l,1,0} & \theta_{l,1,1} & \theta_{l,1,2} & \theta_{l,1,3} & \theta_{l,1,4} \end{bmatrix} \tag{5.4}$$

$$\mathcal{P}(s, \Theta_l) = \begin{bmatrix} s_{0,l}(s) \\ s_{1,l}(s) \end{bmatrix} = \Theta_l \cdot \begin{bmatrix} 1 & s & s^2 & s^3 & s^4 \end{bmatrix}^T. \tag{5.5}$$

The prior smoothness of $\mathcal{C}(s, \Theta_l)$ allows the layer boundaries to be located despite the high speckle noise introduced during the UHROCT imaging process.

To ensure a continuous and robust linear transform for (5.1), the parameters of curve $\mathcal{C}(s, \Theta_l)$ are defined such that the distance between $\mathcal{C}(s, \Theta_0)$ and $\mathcal{C}(s, \Theta_1)$ is always the shortest distance, as opposed to a perpendicular metric, which is less robust due to the effect of noise on local curvature. The curve $\mathcal{C}(s, \Theta_0)$ and the parameters Θ_0 are established such that,

$$\mathcal{C}(s, \Theta_0) \equiv \mathcal{P}(\tau, \Theta_0)\big|_{\tau = \arg\min_\omega \|\mathcal{C}(\omega, \Theta_1) - \mathcal{C}(s, \Theta_0)\|_2}, \tag{5.6}$$

to ensure that the closest points from $\mathcal{C}(s, \Theta_l)$ to the outer boundaries are $\mathcal{C}(s, \Theta_0)$ and $\mathcal{C}(s, \Theta_1)$ for all s and l.

The following list summarizes the properties associated with the corneal shape model defined in equations (5.1) through (5.6):

1. $\hat{\Omega}_{\alpha_\iota}(s)$ is a fourth-order polynomial.

2. The curves $\hat{\Omega}_{\alpha_\iota}(s)$ have C^{inf} geometric continuity [5].
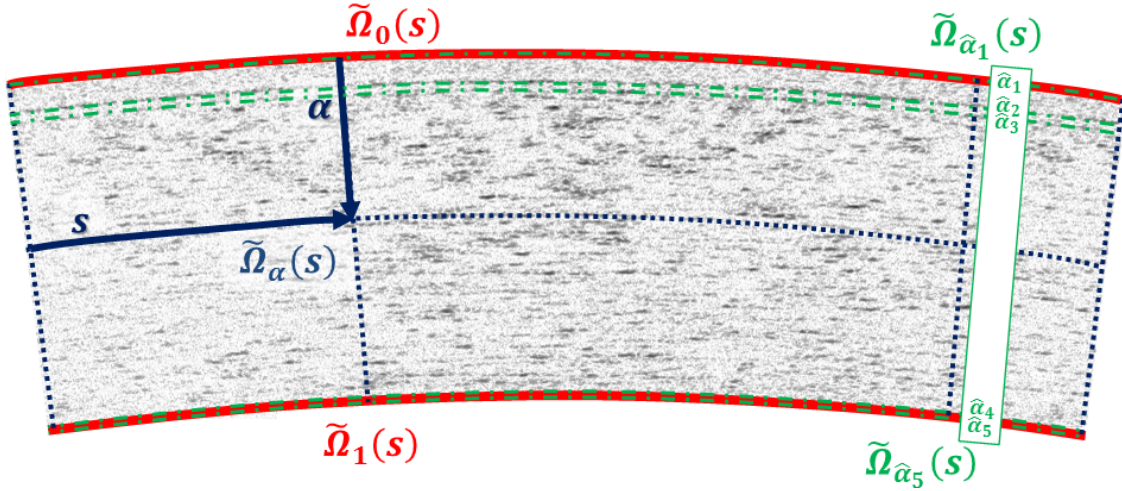
Figure 5.1: The coordinate system used to parameterize each corneal layer boundary illustrates the curve $\Omega_\alpha(s)$, defined in terms of α and s.

3. The Euclidean distance $\left\|\hat{\Omega}_\alpha(s) - \hat{\Omega}_0(s)\right\|_2$ does not decrease as α increases.

What remains is to estimate $\hat{\Theta}_l$ in Section 5.2 and, given the outer boundaries, to find estimates of $\hat{\alpha}_\iota$ in Section 5.3 to locate the inner boundaries.

Since the epithelium and endothelium outer boundaries have the highest contrast, it is comparatively straightforward to extract them using existing algorithms [8]. However, as seen in Fig. 2.5, the inner corneal layer boundaries have lower contrast and cannot be accurately identified by known algorithms, with one example illustrated Fig. 2.10.

However, since the corneal anatomy is layered, with the inner boundaries parallel to the high-contrast outer boundaries, the corneal structure can be modeled. A correspondence model can be established between the upper and lower layers, with the model requiring only a small number of parameters to describe the shape of the inner corneal layer boundaries. Therefore, the corneal layer boundaries are determined through the construction of parameters estimation for such a model such a model and is implemented as an optimization problem.

5.1 Outer boundary segmentation

This section proposes one heuristic approach to segmenting the outer boundaries. The proposed corneal model can be registered to these boundaries allowing the inner boundaries to be segmented, which is the key contribution of automatic corneal layer boundary localization. The parameters $\hat{\Theta}_l$ need to be found such that $\mathcal{C}(s, \hat{\Theta}_l)$ is a good fit for the outermost corneal layer boundaries. Although the outer boundaries have the greatest contrast, image noise and artifacts interfere with segmentation. Since the contrast of the layer boundaries varies over the image, adaptive histogram equalization (AHE) [150] is applied to the OCT image[1]. As a result the boundaries across the entire image have consistent contrast with respect to the surrounding pixel intensities.

In addition, to enhance the arc-like structures of the boundaries, morphological operations [119] are applied. The layer boundaries are enhanced by first applying an open operator with a rectangular 1×7 pixel structuring element to the grey scale tomograms. The angle of the structuring element is iteratively modified so that the operation can preserve arc structures. The process begins with the structuring elements oriented at -30 degrees and ends when the structuring element is at 30 degrees with a change of 3 degrees per iteration. With each morphological iteration, the epithelium and endothelium boundaries become better connected. To reduce the presence of noisy boundaries and produce a more uniform boundary, a large amount of blurring[2] is applied to produce a preprocessed image, presented in Fig. 5.2 (b). The blurring should have the effect to remove all tomogram details while preserving the general corneal shape.

The edges of the cornea are found by applying a horizontal edge detector [99] to the region that represents the cornea. This preserves the predominantly horizontal arc-shape of the cornea while ignoring vertical edges that might be associated with noise. To strengthen the boundaries, the image is dilated with a rectangular structuring operator, as shown in Fig. 5.2 (c). The edge detector and dilation produces sets of points in pixel coordinates, $\Psi_l = \left\{ (s_{0,l,1}, s_{1,l,1}), (s_{0,l,2}, s_{1,l,2}), \ldots \right\}$, for $l = \{0, 1\}$, that correspond to the positions of pixels located on the outer corneal boundaries. All of the detected points are assigned to either the epithelium point set, Ψ_0, or the endothelium point set Ψ_1. Points in the upper half of the image are assigned to the epithelium set, Ψ_0 and points in the lower half of the

---

[1]AHE was applied with the following parameters: the number of tiles is equal to 8×8, the clip limit is 0.01, the number of bins is equal to 256, the full image intensity is used, and a uniform distribution was used.

[2]Gaussian blurring is applied twice sequentially with kernel size of 250 $\mu m \times 37.5$ μm lateral by axial and with sigma equal to 25 μm laterally and 3.75 μm axially. Gaussian blurring is then applied a third time with a size 250 $\mu m \times 37.5$ μm lateral by axial and sigma equal to 50 μm laterally and 7.5 μm axially.

(a) Original

(b) Preprocessed

(c) Edge Detection

(d) Upper/Lower Curves
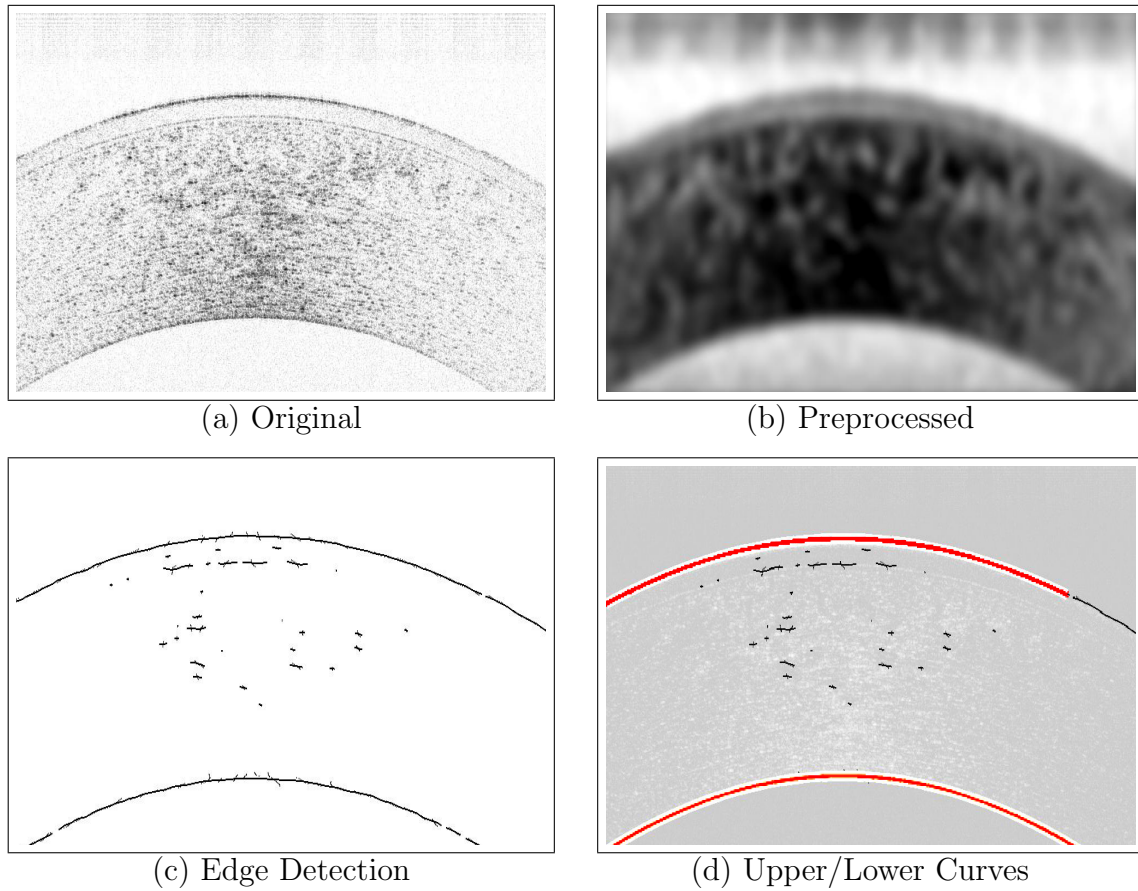
Figure 5.2: Steps involved in the proposed approach to segment the layers of the cornea images. The raw image (a) is preprocessed in (b) so that initial estimates of the upper and lower curves can be calculated. Edge detection (c) is applied to the preprocessed image and (d) the algorithm locates an estimate for the upper, $\hat{\Omega}_0(s)$, and lower, $\hat{\Omega}_1(s)$, curves.

image are assigned to the endothelium set Ψ_1.

5.2 Outer layer parameter estimation

The parameters in Θ_l are estimated so that $\mathcal{C}(s, \hat{\Theta}_l)$ best fits the points in the set Ψ_l. A linear programming model, [104],

$$\hat{\Theta}_l = \arg_{\Theta_l} \min \left(\sum_{\forall k \in \Psi_l} \left(\min_s \|\Omega_l(s) - \Psi_{l,k}\|_2 \right) \right.$$
$$\left. + \sum_{\iota,j} \left(cst_{l,\iota,j} \left(Y_{l,\iota,j,\min} + Y_{l,\iota,j,\max} \right) \right) \right), \tag{5.7}$$

such that:

$$\theta_{l,\iota,j} + MY_{l,\iota,j,\min} > \theta_{l,\iota,j,\min} \qquad\qquad \forall l, \iota, j \tag{5.8}$$
$$\theta_{l,\iota,j} - MY_{l,\iota,j,\max} < \theta_{l,\iota,j,\max} \qquad\qquad \forall l, \iota, j \tag{5.9}$$
$$Y_{l,\iota,j,\min}, Y_{l,\iota,j,\max} \in \{0,1\} \qquad\qquad \forall l, \iota, j, \tag{5.10}$$

is established to ensure that each parameter in Θ_l is between a reasonable minimum and maximum value. The ranges are established with training data consisting of hundreds of segmented healthy tomograms. The indicator variables,

$$Y_{l,\iota,j,\min} \quad \forall \quad l, \iota, j, \tag{5.11}$$

and

$$Y_{l,\iota,j,\max} \quad \forall \quad l, \iota, j, \tag{5.12}$$

are equal to 1 if $\theta_{l,\iota,j}$ is outside the minimum or maximum allowable value, otherwise the indicator variables are equal to zero. The variable M is defined as a sufficiently large constant[3] such that,

$$M \gg \theta_{l,\iota,j} \quad \forall \quad l, \iota, j. \tag{5.13}$$

The constraints (5.8)-(5.10) utilize M, causing $Y_{l,\iota,j}$ to be set to 1 if $\theta_{l,\iota,j}$ exceeds its limits. For every $Y_{l,\iota,j}$ that equals 1, a sufficiently large cost, $cst_{l,\iota,j}$ (5.14), is injected into the cost

---

[3]M was specified here to be $10e10$

function (5.7) whenever the value of the parameter $\theta_{l,\iota,j}$ is not within the range $\theta_{l,\iota,j,\min}$ to $\theta_{l,\iota,j,\max}$.

The optimization function minimizes the distance between each point in Ψ_l and the corresponding curve $\mathcal{C}(s, \hat{\Theta}_l)$ by finding $\hat{\Theta}_l$, the optimal value of Θ_l. The cost function has been designed to improve convergence by penalizing the difference between $\theta_{l,\iota,j}$ and the nearest valid parameter value, $\theta_{l,\iota,j,\min}$ or $\theta_{l,\iota,j,\max}$, when $\theta_{l,\iota,j}$ is outside the valid range. The cost function can be defined as

$$cst_{l,\iota,j} = M \left| \frac{\theta_{l,\iota,j} - \frac{1}{2}\left(\theta_{l,\iota,j,\max} + \theta_{l,\iota,j,\min}\right)}{\theta_{l,\iota,j,\max} - \theta_{l,\iota,j,\min}} \right| \forall l, j. \tag{5.14}$$

Once $\hat{\Theta}_l$ is determined, $\mathcal{C}(s, \hat{\Theta}_l)$ becomes defined allowing the model, $\hat{\Omega}_\alpha(s)$ to be usable. The next step is to accurately locate all five boundaries, $\hat{\Omega}_{\alpha_\iota}(s)$, using the pixel intensities of the original, unprocessed UHROCT tomogram.

5.3 Model guided segmentation

The most significant contribution of this research tool is to present a parameterized layer model that guides the segmentation process allowing low-contrast layer boundaries to be located. The model is registered using the outer layers, $\mathcal{C}(s, \hat{\Theta}_0)$ and $\mathcal{C}(s, \hat{\Theta}_1)$, from Section 5.2, and the layers can be estimated using the model defined by the curves $\hat{\Omega}_{\alpha_\iota}(s)$, from (5.1). Each layer boundary can now be located by determining $\hat{\alpha}_\iota$, the estimate of α_ι. In this instance, the model parameterizes the boundary segmentation into a one-dimensional exhaustive search for $\hat{\alpha}_\iota \forall \iota$ that utilizes an average pixel intensity metric, edge detection, and prior statistics. Ultimately, the actual boundary, $\Omega_{\alpha_\iota}(s)$, is estimated as $\hat{\Omega}_{\hat{\alpha}_\iota}(s)$.

To reduce the possibility of falsely identifying the layer boundary due to imaging noise, the image is implicitly denoised using a spatial weighted average. The average pixel intensity metric, $f(\mathcal{N}, \alpha)$, averages all pixels in the vicinity of the curve $\Omega_{\alpha_\iota}(s) \forall s$ as α varies. The subset of pixels within the surrounding neighborhood is defined to be in the set $\mathcal{N}$. The average pixel intensity metric is a single value that quantifies the intensity of all pixels in $\mathcal{N}$.

The corneal layer boundary is represented by a curve made up of darker intensity pixels surrounded by lighter intensity pixels. A rudimentary edge detector can locate the boundary by taking the average pixel intensity difference between the lighter set and the

darker set of pixels. The set of pixels within the immediate vicinity of $\Omega_{\alpha_\iota}(s)$ $\forall s$ are assigned to the near neighborhood, $\mathcal{N}_{n,\alpha}$. The set of pixels surrounding the potential layer boundary are by definition adjacent to $\mathcal{N}_{n,\alpha}$ and are contained in the far neighborhood set, $\mathcal{N}_{f,\alpha}$.

Pixels near $\Omega_{\alpha_\iota}(s)$ are assigned to either $\mathcal{N}_{n,\alpha}$ or $\mathcal{N}_{f,\alpha}$ based on their distance from $\Omega_{\alpha_\iota}(s)$. To define the distance metric, each pixel in the image, located at $(s_{0,p}, s_{1,p})$ in the image coordinate space, is transformed to the point (α_p, s_p) in the corneal model coordinate space, such that

$$\hat{\Omega}_{\alpha_p}(s_p) = \begin{bmatrix} s_{0,p} \\ s_{1,p} \end{bmatrix}. \tag{5.15}$$

Then the distance from $(s_{0,p}, s_{1,p})$ to $\Omega_\alpha(s)$ is defined as

$$d_\alpha\left(s_{0,p}, s_{1,p}\right) = \alpha_p - \alpha, \tag{5.16}$$

which, from the construction of the model, is the distance from the pixel to the curve projected onto the shortest line segment connecting the outer corneal boundaries. The neighborhood $\mathcal{N}_{n,\alpha}$ contains all the pixels within τ_n pixels of the curve,

$$\mathcal{N}_{n,\alpha} = \left\{ (s_{0,p}, s_{1,p}) \middle| \left| d_\alpha\left(s_{0,p}, s_{1,p}\right) \right| \leq \tau_n \right\}, \tag{5.17}$$

and $\mathcal{N}_{f,\alpha}$ contains pixels surrounding $\mathcal{N}_{n,\alpha}$. The pixels in $\mathcal{N}_{f,\alpha}$ are between τ_n and τ_f pixels away from the curve,

$$\mathcal{N}_{f,\alpha} = \left\{ (s_{0,p}, s_{1,p}) \middle| \tau_n < \left| d_\alpha\left(s_{0,p}, s_{1,p}\right) \right| \leq \tau_f \right\}. \tag{5.18}$$

Due to imaging noise, the average intensity metric, $f(\mathcal{N}, \alpha)$, utilizes a weighting function, W_p, that weights pixels based on their distances from the curve,

$$f(\mathcal{N}, \alpha) = \frac{\int_\mathcal{N} W_p\big(d_\alpha\left(s_{0,p}, s_{1,p}\right)\big) \cdot I(x_p, y_p) d\mathcal{N}}{\int_\mathcal{N} W_p\big(d_\alpha\left(s_{0,p}, s_{1,p}\right)\big) d\mathcal{N}}, \forall (s_{0,p}, s_{1,p}) \in \mathcal{N}. \tag{5.19}$$

The weighting function, W_p, is specific to the noise levels in the tomogram and the thresholds are determined by the spatial imaging scale. The neighbourhoods should be as narrow as possible, setting both τ_n and τ_f to the mean width of the corneal layer boundaries. For noise-free tomograms, W_p should weight pixels further from the curve less than pixels near
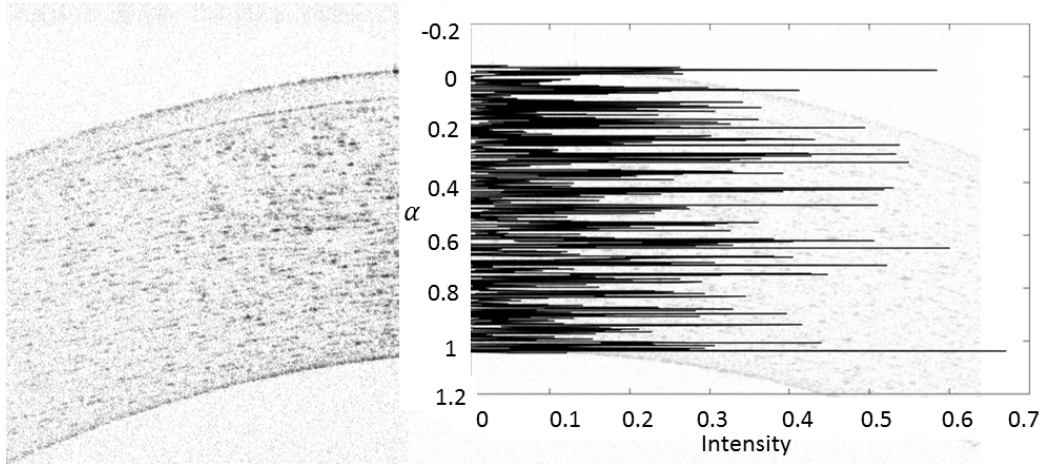
(a) cross-section

(b) average curve intensity for all α

Figure 5.3: (a) The image intensity evaluated at the cross section through the centre of the UHROCT image as α increases from 0 to 1. (b) In contrast with (a), (b) the average intensity along constant α gradients, which is equal to the neighborhood image intensity, $f(\mathcal{N}_\alpha, \alpha)$ with uniform weights and $\tau_n = 0$. The significant minima of $f(\mathcal{N}_\alpha, \alpha)$ correspond to the layer boundaries.

the center of the curve so that the center of the curve generates the most significant edge detection response. However for noisier tomograms, W_p should become more uniform, increasing the amount of blurring thus making boundary location more robust to artifacts, as illustrated in Fig. 5.3.

Since the tomogram is discrete, the average intensity is calculated by sampling the intensities at desired intervals along s and desired distances away from $\Omega_\alpha(s)$. The intensities can be interpolated at the non-integer pixel values obtained when evaluating $\Omega_\alpha(s)$ using a desired interpolation method. For this implementation, n points were uniformly sampled along s, and m points were uniformly sampled in the y direction for every value of s until their distance exceeded τ_n and τ_f. A bilinear interpolation method was applied.

Since the algorithm is attempting to locate layer boundaries, it is expected that the difference between $f(\mathcal{N}_{n,\alpha}, \alpha)$ and $f(\mathcal{N}_{f,\alpha}, \alpha)$, defined as

$$g(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha}) = f(\mathcal{N}_{f,\alpha}, \alpha) - f(\mathcal{N}_{n,\alpha}, \alpha), \tag{5.20}$$

should have a large magnitude to indicate the presence of an edge. Ideally, the five-most prominent values of $g(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha})$ should correspond to $\hat{\alpha}_\iota$ for the five boundaries. However, since there are global contrast variations between the upper and lower pixels in the tomogram, the average contrast varies throughout the image. Consequently, the five-most prominent edges may correspond to false boundaries passing through the center of the tomogram, creating a need for a spatial filter that focuses on statistically likely boundary locations.

Statistical prior knowledge is incorporated into spatial filters that boost the contrast of regions where boundaries are expected and reduce the contrast in between the boundaries. The values of $g(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha})$ are weighted according to a probability density function (PDF), $P(\alpha_\iota)$, corresponding to likely values of α_ι for each boundary. Using segmentation results obtained from hundreds of corneal tomograms, the PDFs for the ι^{th} boundary can be approximated as a Gaussian distribution with a mean of μ_{α_ι} and a standard deviation of σ_{α_ι},

$$P(\alpha_\iota) = \frac{1}{2\pi\sigma_{\alpha_\iota}} \exp\left(-\frac{(\alpha_\iota - \mu_{\alpha_\iota})^2}{2\sigma_{\alpha_\iota}^2}\right). \tag{5.21}$$

Using $P(\alpha_\iota)$ as a spatial weighting, $g(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha})$ is transformed into,

$$h(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha}) = P(\alpha_\iota) \cdot g(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha}), \tag{5.22}$$

for each boundary layer.

72

Finally the location of each corneal layer can be estimated. The five most prominent local minima of $h(\mathcal{N}_{n,\alpha}, \mathcal{N}_{f,\alpha})$, a weighted difference of $f(\mathcal{N}_\alpha, \alpha)$ illustrated in Fig. 5.3 (b), are calculated for $\iota \in \{1, 2, 3, 4, 5\}$. The associated values of α are assigned to each $\hat{\alpha}_\iota$. The resulting curves, $\hat{\Omega}_{\hat{\alpha}_\iota}(s) \ \forall \iota$, represent the location of each of the five layer boundaries, given the tomogram data and the underlying model assumptions. The algorithm corresponding to the entire localization process can be found in Algorithm 1, page 170.

5.4   Evaluation data

The human corneal images for this application are two-dimensional tomograms acquired *in-vivo* from healthy volunteers with a research grade UHROCT system operating in the $1060nm$ wavelength region. Technical details about the system design and performance have been discussed elsewhere [10, 57, 101]. The imaging system provides 3 μm axial and about 15 μm lateral resolution in the human cornea. The optical power of the imaging beam incident on the cornea is limited to $1.5mW$ in accordance with ANSI regulations [30].

Data from a group of twelve healthy human volunteers of different age, gender and ethnicity were recruited for this study and their corneas were imaged with the UHROCT system in compliance with university approved ethics protocols. Two-dimensional corneal UHROCT tomograms were acquired close to the apex of the cornea and have dimensions of 1000×512 pixels (lateral by axial), such that the physical image dimensions are about $5 \, mm \times 1 \, mm$. An example UHROCT image of the human cornea is shown in Fig. 2.5. All corneal layers, except for the endothelium, are clearly visible in the UHROCT tomogram.

The robustness of the automated boundary localization method was validated by segmenting several types of tomogram artefacts acquired during the imaging session. Some tomograms have low optical quality or contain artefacts related to eye motion, saturation of the optical signal due to back reflection from the apex of the cornea, or eyelashes obstructions due to subject blinking. These artefacts are presented in Fig. 5.5.

Because the limited axial resolution of the research grade UHROCT imaging system is $3 \, \mu m$ per pixel, the boundary between Descemet's membrane and the endothelium cannot be resolved. Instead the Descemet's membrane and the endothelium are segmented as a single region, referred to as Descemet's endothelium complex. The data in this document were acquired with an UHROCT system, which allows for non-invasive imaging of a human cornea with 3 μm axial resolution and an acquisition rate of $47,000$ 2D scans per second [33].

Ground truth was established by having a trained technician manually locate each

Table 5.1: Segmentation Experimental Results

Layer	Error			
	Bias [pixels]	Std.Dev [pixels]	Bias [μm]	Std.Dev [μm]
epi	1.32	1.26	3.92	4.38
bow	0.709	10.7	4.59	5.38
str	2.56	5.19	2.86	3.50
des	3.68	4.69	5.94	6.07
end	2.52	3.25	6.52	7.11

corneal boundary. Using a computer, the technician selected between 50 to 100 pixels to represent each inner layer boundary. Due to the labour intensive nature of this process, the technician manually segmented the inner boundaries for 20 randomly selected tomograms and the author, after training, manually generated ground truth for an additional 80 tomograms. Since the outer boundaries had higher contrast and were easier for the technician to segment, the technician provided between 10 and 20 pixels on each outer boundary for all $2,050$ tomograms.

5.5 Results

The proposed algorithm located the epithelium and endothelium boundaries to about 1.32 pixels and 2.52 pixels, respectively, of the manually segmented images for all of the images, with a standard deviation of about 1.3 and 3.2 pixels, respectively. Table 5.1 contains the results in pixels and μm for the other layers. A close-up of the boundary localization results are presented in Fig. 5.4.

The residual error, $r_{l,\iota,auto}$, of the proposed method is quantified as,

$$r_{l,\iota,auto} = s_{1,l,\iota,auto} - s_{1,l,\iota,man},\tag{5.23}$$

where $(s_{0,l,\iota,man}, s_{1,l,\iota,man})$ are the pixel coordinates of the ground truth and $(s_{0,l,\iota,auto}, s_{1,l,\iota,auto})$ are the nearest corresponding points generated using the automated method for each of the layer boundaries. The mean absolute error over all samples, $\bar{e}$,

$$\bar{e} = \sum \frac{1}{n_\iota} \sum_l^{n_\iota} |r_{l,\iota,auto}|,$$

Table 5.2: Error Statistics Measured in Micrometers for each Layer Boundary: mean absolute error ($\bar{e}_\iota$), mean bias ($\bar{b}_\iota$), and the mean standard deviation ($\overline{\sigma}_\iota$)

EIS

layer (i)	$\bar{e}_\iota$	$\bar{b}_\iota$	$\overline{\sigma}_\iota$
epi	5.07	-1.26	9.24
bow	5.76	-3.18	12.42
str	6.42	1.89	15.42
des	6.36	-0.24	14.82
end	4.86	-0.48	8.76

Proposed

layer (i)	$\bar{e}_\iota$	$\bar{b}_\iota$	$\overline{\sigma}_\iota$
epi	3.87	-3.87	**4.86**
bow	12.81	-12.81	**4.56**
str	20.60	-20.60	**4.98**
des	6.00	-6.00	**3.24**
end	6.60	-6.18	**7.62**

Figure 5.4: Layer boundaries are accurately located using the proposed corneal model.

the bias $\bar{b}$, and the standard deviation, $\bar{\sigma}_\iota$, can be used to quantify the accuracy of the obtained layers, and are shown in Table 5.2.

The bias of the EIS algorithm, Table 5.2, is less than that of the automated method due to the fact that the automated algorithm consistently underestimates the location of the resulting corneal boundary. The cause is due to selecting a minimum of $f(\mathcal{N}_\alpha, \alpha)$ instead of selecting a point half way between the minimum and maximum. Since the bias is systematic, the boundary layers are easily corrected by translating the corresponding layer boundaries upward. The standard deviation of the proposed algorithm, is similar for the epithelium layer and notably less for the inner corneal layers, indicating that the fully-automated proposed algorithm is more consistent than a human operator using a semi-automated segmentation process. To further improve the algorithm it may be possible to refine the boundaries given by the proposed method using local optimization techniques. The proposed algorithm, however, is the first published fully-automated algorithm to successfully segment the three inner corneal boundaries [35].

Fig. 5.5 and Fig. 5.6 visually illustrates how well the proposed algorithm identifies each layer boundary compared to EIS. Bowman's membrane and stroma boundary and the Descemet's endothelium complex are indistinguishable to the epithelium and bowman's
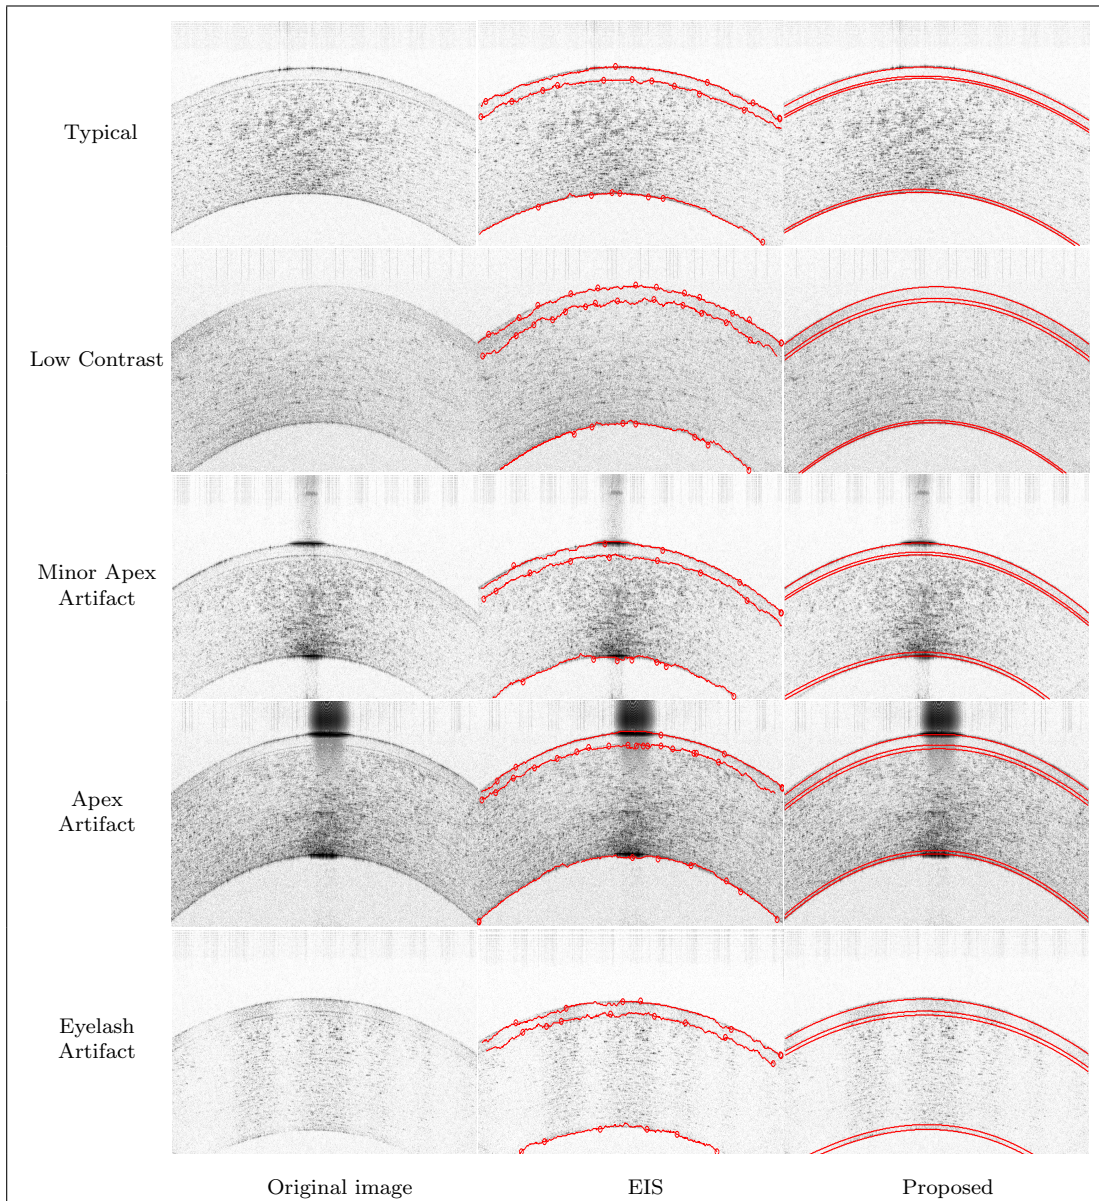
Figure 5.5: Performance of the automated approach compared to EIS for five corneal images. Rows 3 and 4 contains dark regions because the UHROCT images are obtained near the apex of the cornea, and row 5 contains degraded vertical regions due to the obstruction of eye lashes during the imaging process. The automatic method correctly locates all five layer boundaries for each type of image. Although EIS can obtain the general position of the 3 most prominent corneal layer boundaries, inaccuracies of EIS prevent cannot resolve the subtle inner boundaries.

Figure 5.6: Six close-up examples of applying 2D corneal segmentation to OCT images. The model accomplishes exact segmentation, despite significant noise and varying corneal layer location within the image. The prominent imaging artifacts result from the high corneal backscatter due to imaging near the apex.

membrane boundary and the outer endothelium boundary due to the large standard deviation associated with these boundaries, Table 5.2. As seen in rows 3 through 5 of Fig. 5.5, the proposed algorithm performs well despite the presence of imaging artifacts. The results for the 2050 UHROCT data set contain a representative set of UHROCT imaging artifacts.

Chapter 6

Corneal Model

To reduce measurement noise during the imaging process, measurement location uncertainty due to eye movement, and to estimate regions of the cornea that are sparsely sampled, a corneal model is essential as it establishes relationships between neighbouring corneal regions. During the reconstruction process, the corneal model is used to introduce corneal structural, corneal motion, and additional consistency constraints. Since the corneal model establishes corneal layer constraints and a eye motion model, the reader should have an understanding of corneal layers and boundaries, Subsection 2.2.1, their visual appearance within an UHROCT tomogram, Subsection 2.2.2, and an understanding of eye motion, Subsection 2.2.3. Further, while the details of Chapter 5 are supplemental, the structural constraints in this model are derived from the statistics collected in Section 5.5. The reader should also be familiar with Section 2.1 to understand the importance of a subject model and to better understand the model design strategy; Subsection 2.1.2 is essential to this chapter. After the corneal model is created, the resulting model is utilized in Chapter 7. While the details are important for the construction of the model and to satisfy Objective 3, page 45, only high level knowledge of the corneal model is required for understanding reconstruction.

While an infinite number of corneal models can be created, a useful corneal model is one that simplifies the measurement to model correspondence, simplifies the reconstruction process, and allows meaningful data to be extracted. While modelling the cornea at a cellular level allows UHROCT tomograms to be obtained using a UHROCT forward model as one does in reality, the forward model measures coherent light source scattering potential at a specified position, with a complex relationship with the underlying physical characteristics of individual cellular anatomy. Instead, both overall model complexity and forward model complexity can be simplified by modelling scattering potential at any point

within the random field representing the cornea; in this case the forward model can be reduced to a linear measurement correspondence matrix and the intra-relationships between neighbouring elements within the scattering potential 3D field can be modelled. Since one goal of corneal reconstruction is to provide an estimate of the corneal layer thickness in various cross sections, the corneal model should also include a means from which the location of corneal layer boundaries can be easily determined. Further, Fig. 5.3 in Section 5.5 illustrates that scattering potential is coupled with the location of the corneal layer boundaries. The corneal model is designed to decouple these properties, simplifying the structure model parameters and the scattering potential field estimation during the reconstruction process.

Finally, a model of the corneal eye motion is presented as an attempt to bound the error due to measurement location uncertainty. Since the measurement correspondence is part of the forward model, the reconstruction in Chapter 7 estimates the structural and scattering potential field assuming that these values are given. The main use of the eye motion model occurs in Chapter 8 where the measurement location parameters are estimated through an optimization problem designed to minimize reconstruction error.

Synthetic tomograms can then be generated from the corneal model, satisfying Objective 4, page 45. While not essential to the process, they provide easily accessible tomogram data with ground truth for reconstruction and parameter estimation evaluation. Although the model depends on statistics, anatomy, and imaging systems from existing background, the design theory and selective combination of background produce a novel corneal model that facilitates corneal reconstruction and contains an explicit representation of the corneal structure, from which corneal layer thickness measurements can be directly obtained.

Chapter organization

This chapter separates the design of the corneal model into a structural model, a scattering potential field model, and an eye motion model. First, Section 6.1 describes how the 2D layer boundary statistics, presented in Section 5.5, can be expanded to create a 3D surface model for each of the corneal layer boundaries and how the parameters of the surface can be estimated from a random field containing 2D boundary layer statistics collected from all of the available tomograms. Second, Section 6.2 presents Markov random field models that represent the scattering potential relations within each corneal layer, modelling scattering potential as one would model a texture. Third, a simplified eye-motion model is presented in Section 6.3 that quantifies measurement location uncertainty and derives a linear coordinate system transformation allowing measurement location for each pixel

within a tomogram to be repositioned based on eye-motion parameters. Each section also provides synthetic examples of the model for visualization.

6.1 Structural model

The purpose of the structural model is to enforce anatomical corneal structure while performing reconstruction. Since corneal layers are the significant source of structure, modelling theses layers in 3D is the primary objective of this section. The goal of this section is to parameterize each corneal layer boundary in 3D so that it can be calculated efficiently and robustly from the tomogram measurements during reconstruction. For synthesis, the parameters of each layer boundary can be determined through statistical analysis of the 2D corneal layer boundaries described in Section 5.5 and through anatomy described in Subsection 2.2.1. The structural model presented here models healthy human cornea.

This section first presents two different structural models. The first model in Subsection 6.1.2 and extended in Subsection 6.1.3 is useful for creating synthetic cornea because there are a large number of parameters that can be obtained from UHROCT statistics allowing for subjects with greater structural variance. However this model has parameters that are hard to estimate robustly. The second model in Subsection 6.1.4 models the corneal layers as a sphere, reducing the number of parameters required to estimate and allowing the use of existing spherical fitting algorithms. Estimating the spherical model parameters is more robust than the rotating curve parameters thus making it more suitable for reconstruction purposes.

6.1.1 Curvature statistics

Ideally a surface can be constructed from a large dataset of 3D UHROCT containing the location of each corneal layer boundary through manually segmentation. Unfortunately, fully 3D reconstructed corneal data is difficult to obtain and time consuming to manually reconstruct. Further, manually segmented corneal layers within these datasets are non-existent. Instead statistics from individual tomograms combined with expected anatomy can provide corneal curvature information.

The layer boundary model from Chapter 5 models each boundary within a tomogram as a 2D curve. Estimates of each curve,$\hat{\Omega}_{\alpha_\iota}(s)\forall \iota \in [1,2,3,4,5]$, are extracted from the tomogram measurements providing a large collection of corneal curvature measurements.

These measurements can be utilized to convert a 2D curve into a smooth and continuous 3D surface suitable for a structural model.

6.1.2   Rotating curve model

A primitive surface model rotates the first half of the 2D curve, $\hat{\Omega}_{\alpha_\iota}(s)$, oriented in the $\vec{\mathbf{x}}_0$ and $\vec{\mathbf{x}}_1$ plane and centered such that the halfway point of the curve intersects the $\vec{\mathbf{x}}_1$ axis, in 3D space by 360 degrees about $\vec{\mathbf{x}}_1$. Fig. 6.1 (a) illustrates this rotation. The results are smooth and this method models uses the 2D curve allowing it to capture deviations from a spherical corneal model.

6.1.3   Multiple rotating curves model

A slightly better model is to rotate the entire 2D curve using spherical interpolation [120], such that the resulting curve is equal to the first half of $\hat{\Omega}_{\alpha_\iota}(s)$ at 0 degrees and the resulting curve is equal to the second half of $\hat{\Omega}_{\alpha_\iota}(s)$ at 180 degrees, as illustrated in Fig. 6.1 (b). This improvement preserves continuity and surface smoothness but allows asymmetry about the $\vec{\mathbf{x}}_1$ axis.

The spherical interpolation can be extended such that arbitrary curves can be inserted at any desired angle. The surface can be created by interpolating the half curves located at 0 degrees to an independently selected half curve at 90 degrees. Then spherical interpolation continues from 90 to 180 degrees, 180 to -90 degrees, and -90 to 0 degrees. Fig. 6.2 illustrates how surfaces can be generated using this method. While spherical interpolation from multiple curves can be utilized using an infinite number of curves, the model is arbitrarily restricted to the use of two curves, one curve located in the $\vec{\mathbf{x}}_0$ and $\vec{\mathbf{x}}_1$ plane and the other perpendicular curve located in the $\vec{\mathbf{x}}_1$ and the $\vec{\mathbf{x}}_2$ plane. The arbitrary choice is purely based on the amount of effort required to collect corresponding 2D curve statistics from real human tomograms. The result of using this method from two randomly generated curves using real human tomogram statistics is illustrated in Fig. 6.1 (c).

6.1.4   Spherical model

Although the curve rotation models presented in Subsection 6.1.2 and Subsection 6.1.3 are useful for generating synthetic corneal models, fitting data points to multiple rotating curves is difficult and sensitive to noise introduced in sparsely sampled corneal regions. For

(a)

(b)

(c)
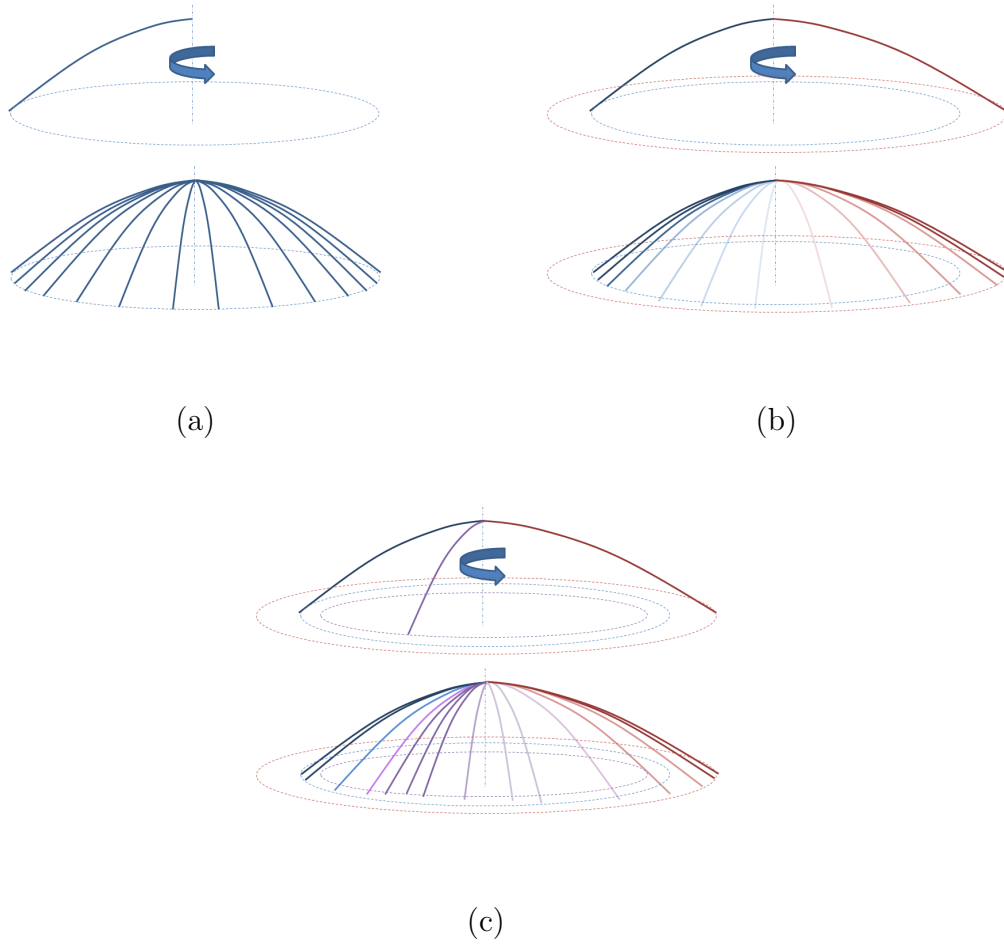
Figure 6.1: (a) shows how a primitive 3D corneal layer boundary surface is created by rotating half of a layer boundary curve. (b) extends (a) by rotating the entire curve from 0 degrees to 180 degrees, using spherical interpolation so that both halves of the curve match the other when rotated 180 degrees. (c) generalizes curve rotation by adding any number of curves to be interpolated.
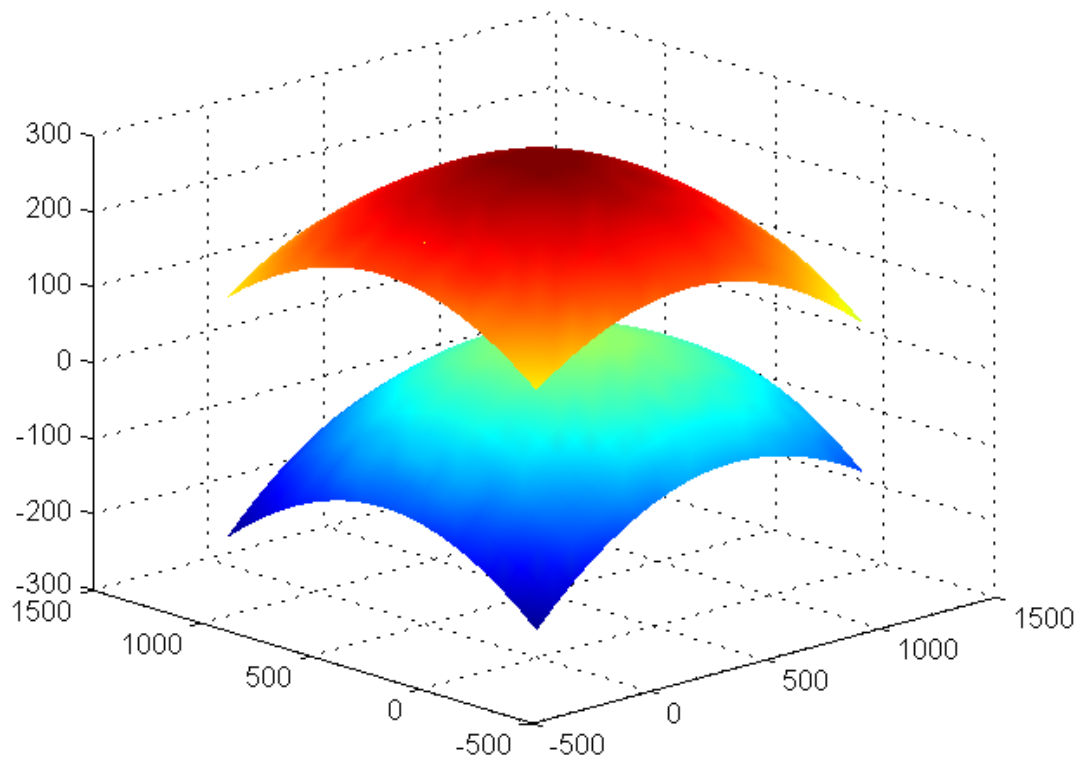
Figure 6.2: The epithelium and endothelium corneal surface layers are generated from spherical interpolation of two perpendicular curves determined through random parameter generation based on curve statistics collected from the UHROCT dataset.

instance, if the layer boundaries from one tomogram contradict a second tomogram the average position from both can be used to reduce measurement noise, or a new curve can be added to the model to perfectly fit both. If the one tomogram is sufficiently far from other tomograms, that tomogram will influence the shape of the curve causing the layer boundaries to be more suspectable to measurement noise. Tomograms in an unsampled corneal region can heavily influence the shape of the cornea, especially if the tomogram's pose has uncertainty.

Instead, a simplistic spherical model with four parameters for each corneal layer is suitable. Since least squares spherical data fitting algorithms already exist [75], there is no need to create a new fitting algorithm from scratch. Further a spherical corneal model is a good approximation of corneal anatomy, shown in Subsection 2.2.1. With only a radius, ρ, and the location of the center, c, for each layer, there are few parameters relative to the number of layer boundary measurements from each tomogram. The parameters can be calculated from only two tomograms sufficiently far apart, for instance on either side of the apex, or two perpendicular tomograms. Additional tomograms improve parameter estimation accuracy.

The parameters for each of the layer boundary radius, ρ_ι and center, $\mathbf{c}_i$, can be represented in the vector Z_1,

$$
Z_1 = \begin{bmatrix} \rho_1 \\ c_{0,1} \\ c_{1,1} \\ c_{2,1} \\ \rho_2 \\ c_{0,2} \\ c_{1,2} \\ c_{2,2} \\ \vdots \\ \rho_5 \\ c_{0,5} \\ c_{1,5} \\ c_{2,5} \end{bmatrix}.
\tag{6.1}
$$

The same layer numbering scheme used in Chapter 5 is also used here. For example, the radius and center of the spherical representation of the epithelium surface is stored as ρ_1 and $(c_{0,1}, c_{1,1}, c_{2,1})$ and the spherical representation of the endothelium layer is stored as ρ_5 and $(c_{0,5}, c_{1,5}, c_{2,5})$. The spherical model is shown in Fig. 6.3, illustrating five corneal layer boundaries.

Figure 6.3: All five corneal layer boundaries, green, blue, magenta, black, and red, are synthesized using parameters sampled from all of the UHROCT boundary location data from Section 5.5.

In isolation the spherical model does not explicitly model local deviations from the sphere. However the location deviations are not removed from the measurements. Instead, when the structure is removed, the location deviations will exist in the relatively structure free scattering potential random field, Z_2. Then the scattering field is reintegrated with the spherical structure model, the local deviations will also be introduced. Thus averaged corneal layer thickness measurements can be taken directly from the spherical model, however more accurate thickness measurements should be taken directly from the combined structural and scattering potential model.

6.2 Scattering potential model

The limited number of human patients and the lack of existing corneal UHROCT datasets with labelled boundaries presents a problem when trying to develop a reconstruction algorithm. While the structure of the corneal layers have been statistically modelled, characterizing scattering potential at a cellular level is more challenging. This section attempts to establish a practical scattering potential model that can characterize the appearance of scattering potential within the corneal layers. While the primary purpose of the model is to estimate scattering potential within a 3D reconstruction, the secondary purpose is to generate simulated UHROCT tomogram ground truth by varying parameters in the structural model, which produces known boundary locations, and to use this scattering potential model to synthesis intra-layer characteristics. Using these models, reconstruction method can be evaluated on a large amount of synthetic data.

While it is possible to extend the scattering potential model beyond appearance, the generation of ground truth for verification and testing is much more challenging than the generation of ground truth for the corneal layer boundaries. For this application, layer boundary segmentation tools were written so that ground truth could be extracted efficiently from human corneal UHROCTs. While having such a tool to model corneal cellular structure is invaluable for other corneal research, the cellular structural modelling falls outside the scope of work for this particular application. Instead, a scattering potential model representing a collection of cells within a corneal layer is more feasible and more robust for the purposes of layer boundary reconstruction than modelling individual cell scattering potential from physics and first principles. The parameters of a regional model are fewer than those required to model the structure of individual cells. The regional model can also be designed using corneal layer statistics that can be segmented using the automated layer boundary localization tool from Chapter 5.

This section is divided into four subsections. The first subsection, Subsection 6.2.1,

decouples the corneal layer structure from scattering potential within an UHROCT tomogram and develops an atlas so that all patient tomograms can be compared and analyzed. The second and third subsections, Subsection 6.2.2 and Subsection 6.2.3, determine how the scattering potential within an UHROCT cross-section can be modelled using a texture model. The fourth subsection, Subsection 6.2.4, shows how the scattering potential from the atlas and be recombined with the structural model so that an UHROCT cross-section can be synthesized.

6.2.1 2D UHROCT atlas

This subsection explores how to develop an UHROCT atlas while simultaneously modelling the scattering potential of the corneal layers from UHROCT tomogram using a random field described in Section 2.1.2.

Since each corneal UHROCT cross-section is imaged at different positions in the cornea for each subject, a consistent cellular density measurement process is necessary to capture the spatial cellular density biases and variances. The intensity of the layer boundaries near the apex of the cornea differ from the intensities further from the anterior pole. Measuring the cellular density as a function of an arbitrary reference frame with basis vectors $\overrightarrow{\mathbf{s}}_{0,i}$ and $\overrightarrow{\mathbf{s}}_{1,i}$ dependent on the imaging perspective and independent of the cornea is simple, but confounds intensity with the position and shape of the cornea within the UHROCT. Instead each UHROCT can be registered to an atlas.

Although each UHROCT cross-section can be registered to any atlas, the form of the atlas can be constrained so that scattering potential can be easily and consistently mapped from the tomogram to the atlas once the cross-section is registered. Since every cross-section must be registered, the atlas should capture characteristics common to all cross-sections. The layer boundaries are convenient common characteristics because they are the most prominent structure in each tomogram and, unlike the location of the anterior pole, are common in all of the tomograms. In addition, the layer boundaries can be automatically located using the automatic layer boundary localization tool, previously discussed in Chapter 5.

Further the use of the layer boundaries reduces the registration problem into a mapping problem; each pixel in T_i can be mapped from the point $(s_{0,p}, s_{1,p})$ to the corresponding point in the corneal model coordinate space using the relationship defined in (5.15). The point $(s_{0,p}, s_{1,p})$ is mapped to (s_p, α_p) and α_p and s_p are linearly transformed into the atlas space. The point $(s_p, \alpha_p) = (0, 0)$ is mapped to $(0, 750)$ and the point $(s_p, \alpha_p) = (1000, 1)$ is mapped to $(1000, 250)$.

Fig. 6.4 represents the atlas corresponding to the registration process and Fig. 6.5 illustrates the registration of a single tomogram. Notice that the air to corneal epithelium boundary and the inner eye to corneal endothelium boundary are straightened during the registration process. This is the result from defining the air to corneal epithelium boundary as $\Omega_{\alpha_p=0}(s)$ and the inner eye to corneal endothelium boundary as $\Omega_{\alpha_p=1}(s)$; since α_p is constant and the registration maps points from $(s_{0,p}, s_{1,p})$ to (α_p, s_p), the epithelium and endothelium boundaries become line segments along $\alpha_p = 0$ and $\alpha_p = 1$, respectively. The (α_p, s_p) space can be quantized so that a bitmap can be used to store the results of registering tomogram T_i.

6.2.2   Simple uniform model

The texture model can be created iteratively by starting with the simplest model and adding complexity until the model is sufficient. The simplest texture is a uniform model that can represent both the vertical and horizonal scattering potential trends using the mean of all tomograms for the same region. Fig. 6.6 represents the mean from all the patients in the dataset. Notice that all five layer boundaries are straightened and appear as dark borders within the composite. The textures in each corneal layer also have interesting characteristics. For instance, the stroma has patches of high scattering potential indicating trends in all of the patients. In contrast, the epithelium, Bowman's membrane, and the Descemet's endothelium complex are largely uniform in intensity.

The trends of the texture can be determined using line integrals through the α direction, shown in Fig. 6.7 and Fig. 6.9. Ignoring the extremities, using vertical line integrals, the horizontal trends create lighter intensity near the edges and darker intensity towards the center. Similarly, using horizontal line integrals, the vertical trends correspond to relatively lower local intensities at each corneal layer boundary, shown in Fig. 6.8. This model can be synthesized by simply extending the results of the vertical and horizonal line integrals in the horizontal and vertical directions, respectively, and taking the mean of the result, which produces the synthetic model shown in Fig. 6.10.

6.2.3   Scattering potential model

Since the synthetic texture model should represent scattering potential of the corneal layers within the tomogram, the previous model from Subsection 6.2.2 provides a good estimate of the average scattering potential, but does not attempt to model local values. Fortunately, a Markov random field with sufficient complexity can parametrically model textures, such

Figure 6.4: The texture registration process adjusts each tomogram to match this atlas. The air-epithelium layer boundary occurs 250 pixels from the top of the image and the inner eye-endothelium layer boundary occurs 250 pixels from the bottom of the image. The other layer boundaries occur approximately at the values specified in the figure but have slight variation in each tomogram. A spatial mapping function specific to each tomogram is required to convert pixel coordinates from the atlas to micrometers.

Figure 6.5: This figure illustrates the result of registering an arbitrary human tomogram to a cornea atlas. The epithelium and endothelium boundaries become straightened during the process. These boundaries are consistently located at the same location for all tomograms.

Figure 6.6: Each tomogram was warped such that the outer corneal boundaries were straightened. Then all of the tomograms were registered and averaged to create the above composite. Vertical and horizontal intensity parameters can be estimated from the composite.

(Vertical Trends)   (Horizontal Trends)

Figure 6.7: Vertical and horizonal trends (blue) are measured using line integrals (red) through the UHROCT composite. The vertical intensities (left-blue) are measured by finding the mean of all pixels in each horizontal line (left-red) while the horizonal pixel intensities (right-blue) are measured by finding the mean of all pixels in each vertical line (right-red).
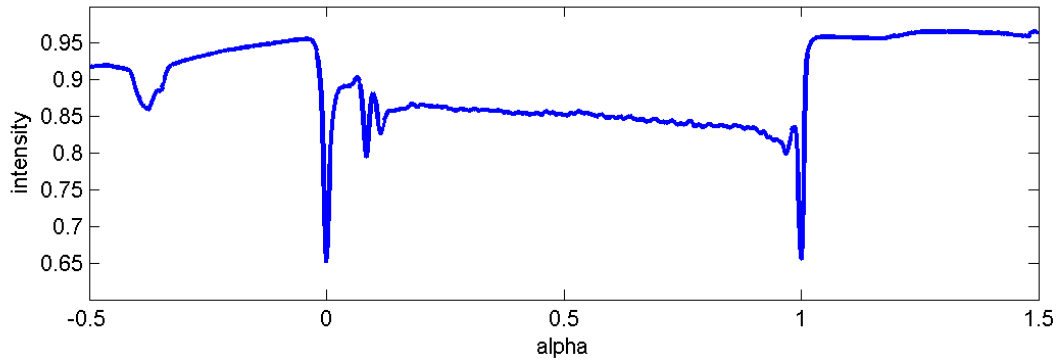
Figure 6.8: The pixel intensity of the composite's vertical trends is measured and illustrated above. Notice that each layer boundary corresponds to a local pixel intensity minima. The air-epithelium layer occurs when α equals zero and the inner eye-endothelium layer boundary occurs when α equals one.

Figure 6.9: The pixel intensity of the composite's horizontal trends is measured and illustrated above. Notice that the apex of the cornea represents a lower pixel intensity than the surrounding edges.

Figure 6.10: Using the intensities for the vertical (Fig. 6.8) and horizonal (Fig. 6.9) trends, a simple synthetic composite can be constructed. Parameters corresponding to the position of layer boundaries can be used to create variations. This simple synthesis does not contain any cellular texture.

as tree bark, cloth, sharp and blurry textures [42]. The uniform model can be combined with a Markov random field to model local trends.

Existing methods to estimate a Markov random field are presented in Subsection 2.1.2. Using these methods, the covariance, P, is calculated for each corneal layer within $I_{i,c}$, the bitmap result of registering tomogram T_i to the corneal atlas. Fig. 6.11 illustrates P associated with the stroma layer for one arbitrary $I_{i,c}$. In addition to covariance for one $I_{i,c}$, the covariance of the entire population is estimated using the dataset of UHROCT tomograms, illustrated in Fig. 6.12. It is important to note that the general structure of P appears to be stationary and periodic since each column of P in Fig. 6.12 can be approximated from the first column, $P_{:,1}$, with a rotation. For example

$$P_{a,b} \approx P_{a-1,b-1}. \tag{6.2}$$

In fact, the mean absolute error satisfies,

$$\frac{1}{n} \sum_{\forall a>1, b>1} \left(\|P_{a,b} - P_{a-1,b-1}\|_2 \right) < 6.15e-005, \tag{6.3}$$

and the mean absolute percentage error satisfies,

$$\frac{1}{n} \sum_{\forall a>1, b>1} \left(\left\| \frac{P_{a,b} - P_{a-1,b-1}}{P_{a,b}} \right\|_2 \right) < 1.18\%. \tag{6.4}$$

The stationary and periodic covariance assumption is supported by the low error terms. The resulting covariance matrix is illustrated in the left image in Fig. 6.12. The dense covariance matrix P can be converted into a kernel $\mathcal{P}$ using material from Subsection 2.1.2.

The statistical model $\mathcal{P}$ is now used to synthesize a cellular density texture. The random field representing the texture is seeded with random initial conditions. Potential seed functions such as Gaussian noise, speckle noise, and salt and pepper noise, and uniform noise can be proposed to simulate the visual appearance of the corneal. Representative samples of the seed images are illustrated in Fig. 6.13.

Using a desired noise function, the random field, N, is populated with the desired seed, such as the sampled CDF seed function. The resulting cellular density texture, Z, is obtained by convolving the square of $\mathcal{P}$ with the randomly seeded field N [42]

$$Z = \mathcal{F}^{-1} \left(\sqrt{\Re \left(\mathcal{F} \mathcal{P} \right)} \cdot \mathcal{F} N \right). \tag{6.5}$$

Figure 6.11: The covariance associated with the stroma corneal layer from an arbitrary tomogram after registration to the corneal atlas is visually represented. The red and yellow colors along the diagonal indicate that pixels near each other have a high correlation while the darker colors further from the diagonal indicate that pixels further from each other have less correlation than the closer pixels.

(covariance from all data)　　　　(periodic and stationary estimate)

Figure 6.12: The covariance matrix on the left is obtained from all subjects, which has less noise than the sample covariance for a single random tomogram illustrated in Fig. 6.11. The covariance matrix on the right is obtained by finding the stationary and periodic properties of the covariance matrix on the left. Notice that each column of the periodic and stationary estimate is a cyclic rotation of the columns next to it, as is expected for stationary and periodic structures.

(Gaussian)   (Salt and Pepper)

(Speckle)   (Uniform)

Figure 6.13: These seed functions, Gaussian, salt and pepper, speckle, and uniform are noise functions that can be used in combination with a Markov model to generate a scattering potential appearance found in a typical UHROCT tomogram. The Gaussian and Speckle noise is smooth while the salt and pepper and uniform noise contains neighbours with larger pixel differences.

The benefit of using a uniform seed is that histogram equalization [96], can applied to Z such that the resulting marginal distribution matches that of the corneal layers. The PDF and CDF for the stroma obtained from a collection of UHROCT tomograms is illustrated in Fig. 6.14 and the inverse CDF is illustrated in Fig. 6.15.

In addition to the statistical model, $\mathcal{P}$, a deterministic model consisting of second-order constraints, $\mathcal{Q}$, derived in Subsection 2.1.2, is illustrated in Fig. 6.16. The $\mathcal{P}$ model estimated from many tomograms has suppressed covariance values due to averaging many covariance matrices together. Both the second-order constraints and the $\mathcal{P}$ estimate from a single tomogram produce reasonable results. The second-order constraints created particulates more similar in size to those in a real tomogram than texture model based on the $\mathcal{P}$ estimate from a single tomogram.

The primary desired trait that is still missing is the horizontal structure that is present in the real tomogram, but lacking from the textures. While $\mathcal{P}$ is calculated from statistics, the deterministic second-order constraint model, $\mathcal{Q}$ can be modified to give more weight to elements in the horizontal, the $\overrightarrow{\mathbf{x}}_0$ and $\overrightarrow{\mathbf{x}}_1$ directions. Fig. 6.17 illustrates the results of the horizontal bias for the deterministic second-order constraint model, $\mathcal{Q}$.

Using $\mathcal{Q}$ the 3D random field, Z_2 can be estimated from a set of existing tomograms or a synthetic cornea can be sampled from the deterministic model using a random seed. The random field estimation is discussed in Section 7.2.2 so that the pixel intensities between tomograms can be estimated from a set of measurements.

6.2.4 Combined atlas and scattering potential model

The Markov random field texture model of the scattering potential and the horizontal and vertical trends obtained from analyzing the UHROCT cross-section atlas represent two different aspects of the UHROCT model. Both the trends and scattering potential model can be combined to create a realistic model of the cross-section. The pixel intensity trends match the format of the corneal atlas and provides a mean to which the scattering potential model can be added. To prevent image saturation and to produce a resulting image that is similar to the human UHROCT cross-sections, the intensities of the vertical and horizontal trends are repeated to create a 3D field the same size as scattering potential field. The two intensity fields are added together so that the resulting intensities closely resembles a human UHROCT cross-section. The combined model containing the Markov random field texture and the structural and illumination biases is illustrated in Fig. 6.18.

The final step is to warp the 3D field of scattering potential, Z_2 to the warped field Z using the structure model from Section 6.1. Points are mapped in 3D in a similar process

(Normalized Histogram)

(Estimated CDF)

Figure 6.14: The PDF is estimated from a histogram of pixel intensities (scaled from zero to one) of stroma regions extracted from UHROCT tomograms. The CDF (red spline) is calculated directly from a cumulative sum of the normalized histogram.

(Inverse CDF)

Figure 6.15: The inverse CDF allows a random uniform density function to be transformed to have the same distribution as the PDF.

to the 2D mapping that used (5.15). Let $\mathbf{x}_{i,2}$ represent a point in N_2 and let $\mathbf{x}_{i,0}$ represent the same point after it is mapped to Z. Also note that when the tomogram data was registered to the corneal atlas, the vertical direction is equal to an affine transformation of the α dimension,

$$x_{2,i,2} = -500\alpha + 750. \tag{6.6}$$

The horizonal and depth components, $x_{0,i,2}$ and $x_{1,i,2}$, are independent. When the structure is reintegrated, the horizonal and depth can remain the same,

$$x_{0,i,0} = x_{0,i,2} \tag{6.7}$$
$$x_{1,i,0} = x_{1,i,2}, \tag{6.8}$$

but the vertical component depends on the vertical position of the epithelium, $x_{2,epi,0}$ and endothelium, $x_{2,end,0}$ layers. Adapting (5.1) for the 3D transform results in

$$x_{2,i,0} = \alpha x_{2,end,0} + (1-\alpha)\, x_{2,epi,0} \tag{6.9}$$
$$x_{2,i,0} = \left(\frac{750 - x_{2,i,2}}{500}\right) x_{2,end,0} + \left(1 - \frac{750 - x_{2,i,2}}{500}\right) x_{2,epi,0}. \tag{6.10}$$

(real tomogram)

(second-order constraint)

(estimated from single tomogram)

(estimated from many tomograms)
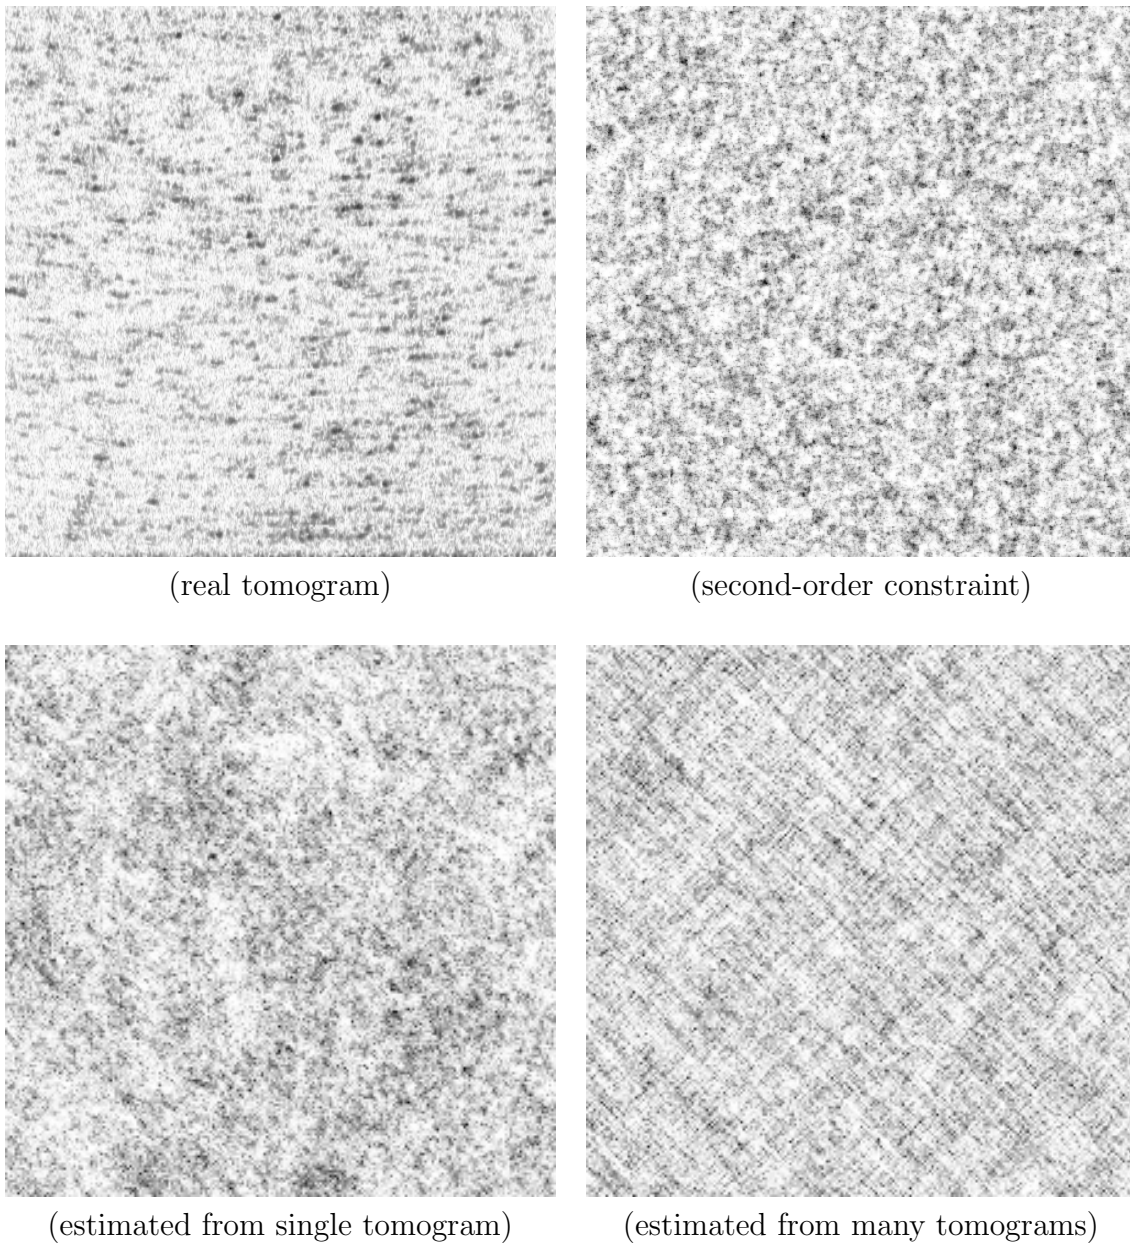
Figure 6.16: The three scattering potential textures above have been sampled using the same uniform noise seed function and different cellular density models. The texture generated from the deterministic second-order constraints is able to produce a texture with large particulates, similar in size to those in a real-tomogram, while the texture generated by estimating $\mathcal{P}$ from a single tomogram does not.

Figure 6.17: Each corneal layer in the above texture is sampled from a second-order deterministic kernel, $\mathcal{Q}$, with horizontal biases. Notice how the resulting texture now has more prominent horizonal noise compared to Fig. 6.16 (second-order constraint).

Then the appropriate structure model is used to determine $x_{2,epi,0}$ and $x_{2,end,0}$ for all the coordinates $x_{0,i,0}$ and $x_{1,i,0}$. While the rotating curve structural model is implemented using Delaunay triangulation [29], the mapping for the spherical model substitutes

$$x_{2,epi,0} = \left(\rho_1{}^2 - \left(x_{0,i,0} - c_{0,1} \right)^2 - \left(x_{1,i,0} - c_{1,1} \right)^2 \right)^{\frac{1}{2}} + c_{2,1} \tag{6.11}$$

$$x_{2,end,0} = \left(\rho_5{}^2 - \left(x_{0,i,0} - c_{0,5} \right)^2 - \left(x_{1,i,0} - c_{1,5} \right)^2 \right)^{\frac{1}{2}} + c_{2,5}. \tag{6.12}$$

into $x_{2,i,0}$.

This formulation allows any number of tomograms to be generated along any specified imaging plane to be sampled from the resulting 3D field, Z. While perpendicular and parallel sampling is exemplified in Fig. 6.19 and Fig. 6.20 respectively, the sampling interface allows the pose of each individual tomogram to be specified facilitating any image plane through the cornea, within the range limitations inherit in the UHROCT dataset.

Arbitrary 2D tomograms can be sampled from the 3D synthetic corneal model, illustrated in Fig. 6.21. The resolution of the tomogram is specified by the operator during sampling and the resolution of the discrete 3D corneal model is specified by the operator

(combined synthetic model)

(real tomogram)

Figure 6.18: A cross-section of the combined cellular texture model (top) is created from a fast fourier transform from a uniform seed generated from a second-order deterministic model for each layer boundary. A real human tomogram (bottom) is provided for comparison.
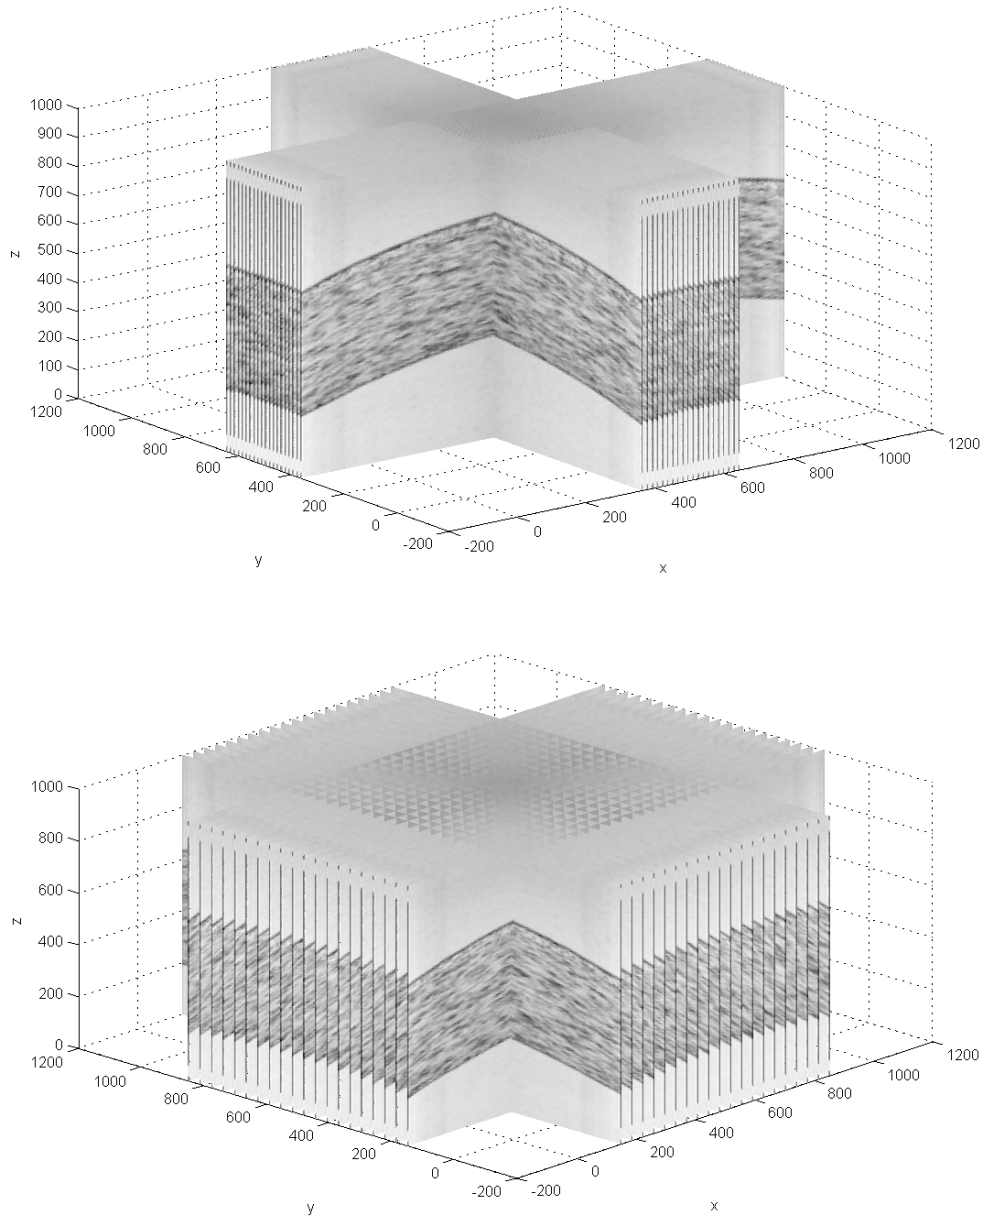
Figure 6.19: 3D synthetic tomogram sets are generated from the corneal model allowing the operate to specify perpendicular sampling focusing on densely sampling the center of the cornea (top) and focusing on a wider perspective of the same subject (bottom).

107

Figure 6.20: 3D synthetic tomogram sets are generated by specifying parallel sampling of the corneal model.

Figure 6.21: A 2D corneal tomogram is sampled from the 3D corneal model. The tomogram sampling resolution for this particular tomogram is specified as 600×1000 pixels, sampled from a synthetic 3D model with resolution $800 \times 800 \times 1000$ pixels.

during synthesis. The image quality of the sampled tomogram is dependent on both the sampling and synthesis resolutions. While high image quality is desired, high quality comes at the cost of increased processing time to create a synthetic model and increased memory usage. For this particular implementation the resolutions were tuned based on available computing hardware such that a 3D corneal model can be synthesized without exceeding $20Gb$ of RAM at any time during synthesis. Initially solid-state disk drive (SSD) paging was implemented to provide an additional $500Gb$ of memory during synthesis, however the computational time required for the RAM to SSD input/output interface was impractical. High-definition 3D models, $2000 \times 2000 \times 2000$ pixels required days to synthesis. Since the implementation was becoming increasingly complex the paging was discontinued and replaced with a highly optimized and memory efficient implementation, allowing textures of up to $10Gb$ given available RAM of $20Gb$.

6.3 Integrated motion model

The reconstruction process can compensate for eye movement during data acquisition using the 3D corneal model synthesized in Section 6.2 by integrating the uncertainty caused by tremors, drift, and microsaccades described in Subsection 2.2.3. The rotations illustrated in Fig. 6.22 gives each point in the UHROCT three degrees of freedom in the tangential direction. This section integrates the eye motion into the corneal model allowing tomograms to be sampled from a synthetic corneal model with representative random eye movement and ultimately allowing reconstruction from real human tomograms that contain such eye motion.

6.3.1 Motivation

Since, as described in Subsection 2.2.2, a single tomogram can be imaged in about 21 to $34ms$, the eye-motion described in Table 2.3 results in a mismatch between the desired location of the last measurement and the actual location of the measurement. Let d be the Euclidean distance between the desired location of the last measurement and the actual location of the measurement. For an average healthy human cornea, eye motion accounts for $d = 26.2~\mu m$ to $d = 153~\mu m$. When viewed from the tomogram, this tomogram is primarily in the tangential direction since, for small angular rotation approximations at the anterior pole,

$$\sin(\theta) \approx \theta \tag{6.13}$$

$$\cos(\theta) \approx 1 - \frac{\theta^2}{2} \tag{6.14}$$

$$\tan(\theta) \approx \theta, \tag{6.15}$$

the tangential direction corresponds to the direction of travel, perpendicular to the radius of the cornea. Since, on average, the Euclidean distance from the anterior pole to the center of rotation is $r_{cornea} = 13.5~mm$ [83], the axial, d_{axial}, and tangential $d_{tangential}$ components of d can be calculated for θ, the angular deviation of the A-scan from the anterior pole, as,

$$d_{axial} = d\sin(\theta) \approx d\theta \tag{6.16}$$

$$d_{tangential} = d\cos(\theta) \approx d\left(1 - \frac{\theta^2}{2}\right). \tag{6.17}$$

(a) sagittal   (b) transverse

Figure 6.22: The tomogram sampling area (orange) represents two tomograms (a) and (b) with the lateral direction centered at the anterior pole. While the image plane for the sagittal tomogram (a) is clearly visible, the image plane for the transverse tomogram (b) is into the page. As the eye rotates at very small angles, the rotation motion results primary in tangential translation, see Fig. 6.23.

For an UHROCT that has resolution of 3×10 μm, (axial $\times$ lateral), and 300×1000 pixels, the angle at the extents, ± 500 pixels from the center, can be calculated as,

$$\tan(\theta) \approx \theta = \frac{10 \frac{\mu m}{pixel} \cdot 500 pixels}{r_{cornea} \ \mu m} = \frac{5000}{13500} = 0.37 \tag{6.18}$$

corresponding to d_{axial} and $d_{tangential}$, 56.7 and 142 μm, given $d = 153$ μm. Fig. 6.23 illustrates d_{axial} and $d_{tangential}$ at several positions within a tomogram centered at the anterior pole.

As presented in Fig. 6.23, d for a single tomogram is reasonable, and on average, does not exceed 153 μm. However, as the number of tomograms increases, d also increases. In UHROCT imaging, the patient is asked to stare at a target. If a target was not present, d could increase until bounded by physical rotation limits or, more realistically, by environment and patient specific neurological impulses. However, the target exists, and as d becomes too large, theories suggest that the microsaccades refocus the eye towards the

Error (red regions) due to drift after 0.034 s			
region	axial (μm)	lateral (μm)	area (mm^2)
(a)	56.7	142	0.0081
(b)	25.5	151	0.0039
(c)	0.00	153	0.0000
(d)	59.2	142	0.0084
(e)	26.6	151	0.0040

Figure 6.23: When eye motion occurs during UHROCT sampling measurement location error (red region) occurs. The error (c) at the anterior pole is a horizonal line since the only allowable motion is lateral due to the anatomy of the eye and 3D rotational axes. As the tomogram sampling moves further from the anterior pole, from (c) through (b) to (a), the eye motion allows the volume containing the measurement location error increase; note that the lateral direction also expands into and out of the page and only the cross-section of the volume is shown in this figure. Note that (a) and (d) have the same pixel location, not the same axial location. This figure is to scale assuming a spherical corneal assumption, an epithelium radius of 7.5 mm [61], and a central corneal thickness of 581 μm [8].

target every 0.23 Hz to 2.5 Hz [78], or every 0.40 to 4.35 seconds. During 0.40 seconds, between 10 to 20 UHROCT tomograms can be obtained, but the anterior pole might move as much as 1740 μm. Fig. 6.24 presents a visualization of d after 0.40 seconds of sequential UHROCT sampling and Fig. 6.25 compares d for after 0.034, 0.40, and 1.6 seconds, respectively.

At the upper limit of the microsaccades period, 4.35 seconds, the anterior pole might move as much as $18,800$ μm or 54.4 deg, which is nearly the entire vertical field of view, and should be discarded by the operator if all corneal layers would no longer be visible in the tomogram. Since corneal regions appear very similar, however, the UHROCT tomogram at the desired measurement location can be mistaken for many other corneal regions resulting in great difficulty and location uncertainty when attempting to measure a desired location. The rest of this section parameterizes this location uncertainty so that the actual measured position can be estimated based on all available tomograms. Note that while the last tomogram in a sequence may have a large d relative to the first tomogram in the same sequence, the d within the same tomogram is relatively small since the sampling period for a single tomogram is only 0.021 to 0.034 seconds; the location of measurements within the same tomogram are relatively consistent.

6.3.2   Motion model

The large amount of measurement uncertainty during sequential tomogram imaging due to eye motion, discussed in Subsection 6.3.1, justifies the need to model the uncertainty of the tomogram pose, $\mathbf{\Theta}_i$, previously defined in Subsection 2.1.2. This section establishes consistent relationships between the elements of $\mathbf{\Theta}_i$ and the eye motion.

Modelling only the eye motion, the vector connecting the position of the corneal axes of rotation to the UHROCT imaging system remains constant. Hence, establishing the origin of the coordinate system at the corneal axes of rotation ensures that the coordinate system does not translate relative to the UHROCT system. Further, a spherical coordinate system can represent three rotational degrees of freedom, in the form of pitch, roll, and yaw Euler angles $(\varphi, \psi, \vartheta)$ required to model human eye motion at any given point in time. The radial distance (ρ) is determined as the distance from the origin of the spherical coordinate system to the desired sampling point, $\mathbf{x}_{i,k}$, specified by the UHROCT imaging

Error (red regions) due to drift after $0.40\ s$			
region	axial (μm)	lateral (μm)	area (mm^2)
(a)	644	1620	1.0433
(b)	290	1720	0.4988
(c)	0	1740	0.0000
(d)	673	1610	1.0835
(e)	303	1710	0.5181

Figure 6.24: Similar to Fig. 6.23, which illustrates accumulated error after sampling for $0.034\ s$, after sampling for $0.4\ s$, or about 10 to 20 sequential UHROCT tomograms, eye motion can cause the following d_{axial} and $d_{tangential}$ error at the epithelium and endothelium layers.

Error due to drift in region (b)				
region	time (s)	axial (μm)	lateral (μm)	area (mm^2)
(b.i)	0.034	25.5	151	0.0039
(b.ii)	0.400	290	1720	0.4988
(b.iii)	1.600	1150	6840	7.8730

Figure 6.25: The amount of drift error is represented as cross-sections, to scale, of the volumes containing the desired measurement are illustrated for point (b), for 0.034 s (b.i) from Fig. 6.23, 0.40 s (b.ii) from Fig. 6.24, and 1.60 s (b.iii). After 1.60 s, or the time required to sample about 40 to 80 sequentially sampled tomograms, the error is significantly larger, 7.57 mm, and larger than the radius of the cornea.

system. The corresponding homogenous rotation matrix, H_{rot},

$$H_{rot0}(\psi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\psi) & -\sin(\psi) & 0 \\ 0 & \sin(\psi) & \cos(\psi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{6.19}$$

$$H_{rot1}(\varphi) = \begin{bmatrix} \cos(\varphi) & 0 & -\sin(\varphi) & 0 \\ 0 & 1 & 0 & 0 \\ \sin(\varphi) & 0 & \cos(\varphi) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{6.20}$$

$$H_{rot2}(\vartheta) = \begin{bmatrix} \cos(\vartheta) & \sin(\vartheta) & 0 & 0 \\ -\sin(\vartheta) & \cos(\vartheta) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{6.21}$$

$$H_{rot}(\vartheta, \varphi, \psi) = H_{rot0}(\psi)H_{rot1}(\varphi)H_{rot2}(\vartheta) \tag{6.22}$$

which can rotate a point centered at the spherical origin about the spherical origin given $(\vartheta, \varphi, \psi)$.

The pose parameters in $\boldsymbol{\Theta}_i$ are specified by the UHROCT imaging system but instead of sampling at the desired location $\mathbf{x}_{i,k}$,

$$T_i(\mathbf{s}_{i,k}) = f\big(Z, \mathbf{x}_{i,k}\big), \tag{4.20 on page 56}$$

the desired sampling point $\mathbf{x}_{i,k}$ of tomogram i and pixel k is rotated about the spherical origin to determine the actual sampling point, $\mathbf{x}_{a,i,k}$,

$$\mathbf{x}_{a,i,k} = H_{rot}(\vartheta, \varphi, \psi)\mathbf{x}_{i,k}. \tag{6.23}$$

The tomogram sampling function, (4.20), becomes

$$T_i(\mathbf{s}_{i,k}) = f\big(Z, \mathbf{x}_{a,i,k}\big), \tag{6.24}$$

or by substitution,

$$T_i(\mathbf{s}_{i,k}) = f\big(Z, H_{rot}(\vartheta, \varphi, \psi)\mathbf{x}_{i,k}\big). \tag{6.25}$$

The motion parameters ψ and φ are determined from the tremors, drifts, and microsaccades. Tremors, despite actually being aperiodic, are modelled using sinusoid at $90\ Hz$ and an amplitude of $0.0218\ deg$ for simplicity, drifts are modelled as random walk over time

116

with a speed of 0.41 deg per second and random direction calculated from $\theta_{rand,d}$. The random walk is discretely applied every Δt, where Δt is equal to the UHROCT sampling period,

$$\psi(t + \Delta t) = \psi(t) + 0.0072\Delta t \cos(\theta_{rand,d}) + 2.20 \times 10^{-4}\Delta t \cos(180\pi t) \tag{6.26}$$
$$\varphi(t + \Delta t) = \varphi(t) + 0.0072\Delta t \sin(\theta_{rand,d}) + 2.20 \times 10^{-4}\Delta t \cos(180\pi t), \tag{6.27}$$

and microsaccades are modelled as periodic jumps at 2.4 Hz of random distance, $d_{rand,m}$, between 2.06 deg and 9.64 deg in a randomly improving direction, where $\theta_{rand,m}$ is within 90 deg of the direction of steepest descent for $[\psi(t), \varphi(t)]$. Microsaccades update $\psi(t)$ and $\varphi(t)$ every 0.417 seconds during the simulation,

$$\psi(t) = \psi(t) + d_{rand,m} \sin(\theta_{rand,m}) \tag{6.28}$$
$$\varphi(t) = \varphi(t) + d_{rand,m} \cos(\theta_{rand,m}). \tag{6.29}$$

The ψ is modelled to only be affected by tremors,

$$\vartheta(t + \Delta t) = \vartheta(t) + 2.20 \times 10^{-4}\Delta t \cos(180\pi t). \tag{6.30}$$

Note that the termor is arbitrarily applied as $2.20 \times 10^{-4}\Delta t \sin(180\pi t)$ equally to each rotational axis.

The resulting simulated random eye motion is illustrated in Fig. 6.26, which can be compared with actual measured eye motion in Fig. 2.8. While the simulation correctly models the range of motion when compared to real measurements, several nuances are not modelled, and are not present in the simulation. For instance, the exact strategy of how microsaccades correct drift, i.e. the angle of correct, the amount of correct, and the circumstances affecting the correction period, are intentionally not modelled. The drift was modelled after a random walk, and could instead be modelled as a second-order random walk to reduce the variability of the randomly chosen direction. Further, while the tremor is simulated, the effects are negligible compared to drift and microsaccades and can be removed from the motion model. While the model presented here can incorporate these elements in the future, the eye motion model correctly accounts for the amount of eye motion that can occur while obtaining an A-scan and is utilized to add measurement uncertainty into the forward model and to illustrate reasonable motion limits .

Figure 6.26: The green curves represent the random drift while the black lines represent microsaccades. When compared to Fig. 2.8, the direction of simulated microsaccades does not seem to drift correct toward the desired target of $(0,0)$ as well as the measured eye motion. Further the simulated drift is modelled as a random walk, where the real drift motion appears to follow a smoother curve. The patch has a width and height of 50 μm and is located on the cornea.

Chapter 7

Corneal Reconstruction

Once the measurements are obtained, the corneal model is created, and the forward model is defined, corneal reconstruction can be performed to estimate states and parameters of the corneal model from the measurements. This chapter implements the reconstruction process of Section 2.1, satisfying Objective 5, page 45, using the corneal model and modified forward model defined in Chapter 6, for which the reader should be familiar. In addition to using the reconstruction process described in this chapter to produce a 3D corneal volume and 3D surfaces for each corneal layer boundary, Chapter 8 uses the reconstruction process extensively when estimating parameters associated with the forward model.

Using real patient tomograms and synthetic corneal tomograms, the structural and scattering potential reconstruction methods are presented in this chapter. A preprocessing step is applied to the tomogram that decouples the corneal structure from the UHROCT measured scattering potential in the same way decoupling was applied in the corneal model; preprocessing locates the layer boundaries using Chapter 5 and each tomogram is registered to the UHROCT tomogram atlas defined in Subsection 6.2.1. Reconstruction uses the scattering potential that is mapped into the tomogram atlas and the boundary layer location information to estimate corneal model states. While parameter estimation by itself is not new, the corneal model state estimation and corresponding reconstruction is novel for the UHROCT corneal application area. The ability to measure the corneal thickness through an arbitrary cross-section is the primary contribution resulting from the reconstruction, satisfying Objective 7, page 45.

Chapter organization

This chapter is split into three sections. Section 7.1 presents the tomogram sets from healthy patients and synthetic cornea that are reconstructed and measured in the following sections. Section 7.2 presents all of the parameter estimation from tomograms registered to the tomogram atlas. Then the parameters are used to create a 3D volumetric representation of the corneal layers and a second 3D volumetric representation of the scattering potential. Finally, Section 7.3 presents the reconstruction results using the evaluation data presented in Section 7.1.

7.1 Evaluation data

The UHROCT data obtained from human subjects are described in Section 5.4. For 3D reconstruction, three healthy human subjects were imaged. Ground truth for the data was determined through manual segmentation and reconstruction for use in reconstruction evaluation. The data is summarized in Table 7.1. Although data from four human subjects were collected, initial problems with the imaging procedure prevented the use of some data because the cornea moved out of the imaging plane as one of the subjects was being scanned.

In addition, while over 100 synthetic subjects were created from the corneal model described in Chapter 6 to refine the reconstruction process, 15 synthetic subject sets were randomly generated to be used for reconstruction tests. Their characteristics are summarized in Table 7.2. The corneal layer boundary structure for these subjects was generated using the corneal structure model, in Section 6.1, with randomly selected corneal curvature based on statistics collected from the human dataset, Section 5.4. These synthetic test subjects were synthesized using the scattering potential model, in Section 6.2, and warped to produce scattering potential sandwiched between the corneal structure. Further, three different tomogram sampling schemes were utilized: densely centered sets of sagittal and transverse sampled tomograms (Fig. 6.19 (top)), evenly spaced sets of sagittal and transverse sampled tomograms (Fig. 6.19 (bottom)), and evenly spaced sagittal sampled tomograms, (Fig. 6.20). Fig. 7.1 illustrates the sagittal and transverse directions for corneal sampling.

Figure 7.1: Light enters through the epithelium and exits through the endothelium layers illustrated in blue. The cornea is sampled in the sagittal (pink) and transverse (green) directions.

Table 7.1: Three human corneas, each sampled twice, were manually segmented to produce ground truth for 6 tomogram sets.

Set	Subject	Num. Samples	Sampling Scheme	Sampling Spacing
1	1	44	sagittal	evenly spaced
2	1	44	transverse	evenly spaced
3	2	175	sagittal	evenly spaced
4	2	50	transverse	evenly spaced
5	3	92	sagittal	evenly spaced
6	3	40	transverse	evenly spaced

Table 7.2: Five synthetic corneas were each sampled three different ways to produce 15 tomogram sets used to evaluate reconstruction.

Set	Subject	Num. Samples	Sampling Scheme	Sampling Spacing
1	1	40	sagittal/transverse	dense center
2	2	40	sagittal/transverse	dense center
3	3	40	sagittal/transverse	dense center
4	4	40	sagittal/transverse	dense center
5	5	40	sagittal/transverse	dense center
6	1	40	sagittal/transverse	evenly spaced
7	2	40	sagittal/transverse	evenly spaced
8	3	40	sagittal/transverse	evenly spaced
9	4	40	sagittal/transverse	evenly spaced
10	5	40	sagittal/transverse	evenly spaced
11	1	40	sagittal	evenly spaced
12	2	40	sagittal	evenly spaced
13	3	40	sagittal	evenly spaced
14	4	40	sagittal	evenly spaced
15	5	40	sagittal	evenly spaced

7.2 Decoupled reconstruction

This section utilizes the decoupled corneal model designed in Chapter 6 to reconstruct a cornea. Subsection 7.2.1 preprocesses the measurements $\mathbf{m}$, extracts boundary information using the automated boundary localization method, in Chapter 5, and estimates, Z_1, the corneal structural parameters from the measurements. Subsection 7.2.2 performs weighted least-squares to estimate the Markov random field, Z_2, representing the scattering potential at each point within the reconstruction volume. After the decoupled states Z_1 and Z_2 are estimated independently, they are combined in Subsection 7.2.3 to produce a full corneal reconstruction, Z, a Markov random field containing scattering potential with structural constraints.

7.2.1 Structure

Referring back to Subsection 3.2.2, the corneal measurements are related to the structural states through

$$\wp_1\left(\mathbf{m}\right) = C_1 Z_1. \qquad\qquad \text{(3.7 on page 49)}$$

The function $\wp_1(\mathbf{m})$ is the result of the automatic corneal layer boundary localization tool applied to all of the tomograms. Let Υ represent the location in $\Re^3$ of each measurement on each corneal layer boundary such that,

$$\Upsilon = \wp_1(\mathbf{m}). \qquad\qquad (7.1)$$

Further, each layer boundary is isolated by notating Υ_ι as a subset of Υ containing all of the measurement locations for corneal boundary layer ι,

$$\Upsilon = \{\Upsilon_1, \Upsilon_2, \Upsilon_3, \Upsilon_4, \Upsilon_5\}. \qquad\qquad (7.2)$$

Then, using the spherical structure model designed in Section 6.1 and a least-squares spherical fitting algorithm [75], Algorithm 2, page 171, the radius, ρ_ι, and center, c_ι, of a sphere that fits all of the data in Υ_ι can be estimated for each corneal layer boundary.

Since there are few parameters to estimate and a large number of samples, this process is robust to imaging noise. Since structure is based on a model of corneal anatomy, the boundary extrapolation in unsampled corneal regions is also robust. Fig. 7.2 illustrates the result of the structure fitting process. Fig. 7.2 (a) uses color to identify all of the measurement locations for each layer boundary after they are discovered using $\wp_1$.

Fig. 7.2 (b) shows how the corneal structure is extrapolated and how well it lines up with the measurements.

It should be noted, that at this stage, corneal layer thickness measurements can be obtained using the spherical corneal model. However, this structural model does not take any local deviations into account. Consequently, a more accurate measurement can be obtained when the scattering potential is combined with the structural model in Section 7.2.3.

7.2.2 Scattering potential

The relationship between the measurement scattering potential, $\mathbf{m}$, and the states, Z_2 from Section 6.2, is reproduced from Subsection 3.2.2,

$$\wp_2\left(\mathbf{m}\right) = C_2 Z_2. \qquad \text{(3.8 on page 49)}$$

Instead of using the locations of the corneal boundaries, the $\wp_2(\mathbf{m})$ becomes a set of measurements corresponding to straightened corneal tomograms; each point in the tomogram is transformed from (s_0, s_1) to (α, s) and then mapped into $\mathbf{x} \in \Re^3$ using the transformation derived in Section 4.1. While the pose of the tomogram and the parameter s have been transformed into x_0 and x_1, the scattering potential field was designed to decouple structure from scattering potential, forcing x_2 to correspond directly to α.

The decoupled scattering potential states, Z_2, can be estimated using a weighted least-squares formulation [42], μ_2, the mean of $\wp_2(\mathbf{m})$, and the previously designed deterministic constraints, Q,

$$\hat{Z}_2 = \left(C_2{}^T R^{-1} C_2 + Q\right)^{-1} \left(C_2{}^T R^{-1} \left(\wp_2\left(\mathbf{m}\right) - \mu_2\right)\right) + \mu_2. \qquad (7.3)$$

The mean, μ_2, is first removed so that any scaling factors due to R or Q scale only the differences in measurements and not their overall intensity. The mean is added back into $\hat{Z}_2$ after the scaling process. Pseudocode, Algorithm 3, page 171, solves for $\hat{Z}_2$ based on (7.3).

The resulting random field, $\hat{Z}_2$, is illustrated in Fig. 7.3 for an evenly spaced synthetic cornea and in Fig. 7.4 for a synthetic cornea with dense tomogram sampling near the apex of the cornea. When comparing these two figures, notice that Fig. 7.3 (d) has more noise near the left and right edges of the image compared to Fig. 7.4 (d). The proximity of measurements near the image edges in Fig. 7.3 have greater influence than measurements

(a)

(b)

Figure 7.2: (a) The parameters of the corneal structure model are estimated using parameter estimation from a set of tomograms. (b) Reconstruction of unsampled regions is achieved through the use of the corneal model, incorporating anatomy. The corneal layer boundaries are highlighted through color: outer-epithelium (green), epithelium-Bowman's layer (blue) Bowman's layer-stroma (magenta), stroma-Descemet's membrane (black), and endothelium-outer (red).

125

further away, as in Fig. 7.4. In unsampled corneal regions, the nearest measurements have the greatest impact on scattering potential estimation while the furthest measurements have less impact. In these regions, there are few measurements to perturb the estimate from a mean, resulting in a tendency to have more uniform scattering potential estimates. While arbitrary scattering potential noise can be simulated based on other regions, the estimate from (7.3) is closer to the mean of scattering potential in that region, and therefore more appropriate.

The success of the reconstruction can be seen by comparing Fig. 7.4 (b) to Fig. 7.4 (d). Using every tomogram available, the scattering potential in Fig. 7.4 (d) has been reconstructed where no previous data existed.

7.2.3 Combined structure and scattering potential

The combined reconstruction reintegrates the previously estimated structure and scatter potential to calculate $\hat{Z}$ from $\hat{Z}_1$ and $\hat{Z}_2$ using the same process as described in Subsection 6.2.4, but reiterated here using a reconstruction specific context. Pseudocode for a reconstruction implementation is presented in Algorithm 4, page 172.When the tomograms were straightened in Section 7.2.2, each tomogram was warped such that epithelium boundary layer and the endothelium boundary layer were straightened and the measurements between were transformed from (s_0, s_1) to (α, s). Then each point is positioned from (α, s) to $(x_0, x_1, x_2)_2$ in the reconstructed volume $\hat{Z}_2$. Integrating the structure into the scattering potential field maps $(x_0, x_1, x_2)_2$ to (α, s).

The position of a point in Z_2 can be mapped to Z using three pieces of information. First, x_2 in the context of Z_2 corresponds directly to α. Second α relates the position at (α, s) to a relative distance between the epithelium and endothelium corneal boundary layers. Third, the position of the epithelium and endothelium corneal boundary layers are calculated directly from $\hat{Z}_1$. Essentially this results in $\hat{Z}_2$ becoming warped and constrained between the estimated boundary layers.

Fig. 7.5 illustrates the result of warping the scattering potential random field when combined with the estimated corneal structure. Fig. 7.5 is the complete reconstruction from a synthetic corneal model with densely centered tomogram sampling.

(a)   (b)

(c)   (d)
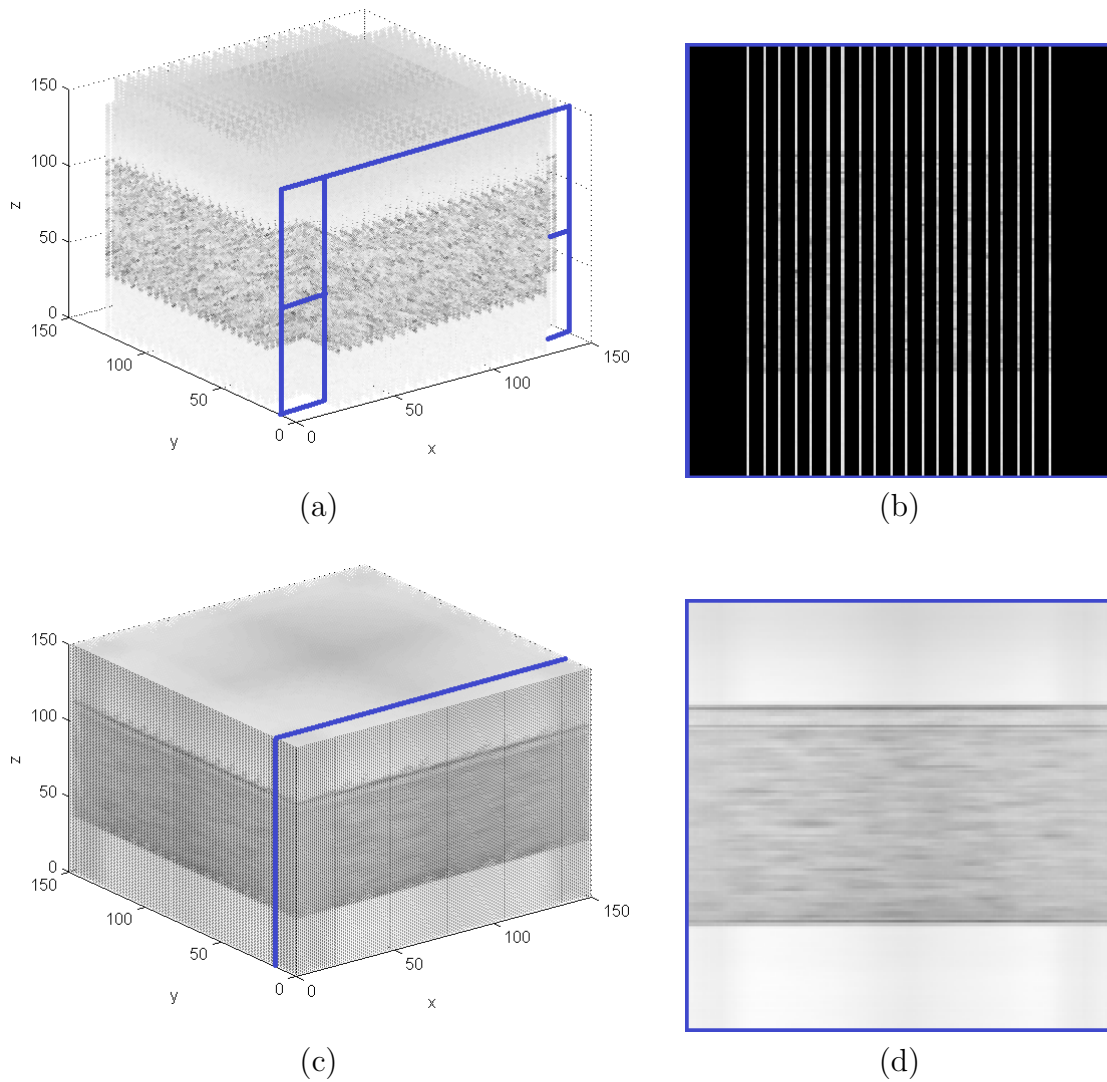
Figure 7.3: Synthetic data sampled evenly using a perpendular and parallel tomogram sampling scheme is illustrated in (a).(a) is sampled at the cross-section in blue to produce (b). The measurements from (a) are reconstructed in (c) and similarly (c) is sampled at the same cross-section to produce (d). Notice in (d) that the left and right edges are slightly more uniform than the center regions. While the boundary layers in (b) can hardly be seen, they can easily be seen in (d), the reconstruction.

Figure 7.4: Similar to Fig. 7.3, (a) is sampled at the cross-section in blue to produce (b) and used to reconstruct (c). Similarly (c) is sampled at the same cross-section to produce (d). However, in this case, the synthetic subject is densely sampled near the apex. Notice that the left and right edges in (d) are much more uniform than those in Fig. 7.3 (d) due to the proximity of measurements to the estimated states.

(a)

(b)

Figure 7.5: (a) The structure and scattering potential are combined to produce a full 3D reconstruction for the cornea. Since, the structure has been previously calculated the corneal boundary layers can be quickly labelled (b). Note that for illustrative purposes some pixels were removed from the 3D visualization so that the corneal structure can be seen. As a result many reconstruction details, such as texturing visible in Fig. 7.2.3 (d) are not visible above. Further, the circular pattern is also an artefact created when attempting 3D visualization. The corneal layer boundaries are highlighted using color in (b).

7.3 Results

The reconstruction results for the synthetic and human subjects are presented in the following two subsections.

7.3.1 Synthetic corneal model

Each of the five synthetic subjects were reconstructed three times using different tomogram sampling schemes: for each synthetic subject, 1) 20 samples in the sagittal direction and 20 samples in the transverse direction with a dense concentration near the center of the cornea, 2) 20 samples in the sagittal direction and 20 samples in transverse direction with an even spread concentration, and 3) 40 samples in the sagittal direction with an even spread concentration.

Table 7.3 and Table 7.4 detail the mean distance from each boundary layer to the corresponding point on the corneal structure model. As the tables and the previously illustrated Fig. 7.2 suggest, the reconstruction error associated with the structural model is less than 0.17 μm on average over all of the layers. This indicates that the reconstruction error for corneal structure model is significantly less than the width of the thinnest corneal layer, 5 μm.

Similarly, the reconstruction error associated with the scattering potential was negligible. Since the synthetic models were derived exactly from the scattering potential model, the scattering potential, scaled from 0 to 1, for all three subjects is summarized in Table 7.5. The difference between tomogram scattering potential and reconstructed scattering potential was insignificant, less than 0.0178 for 50% of all pixels for all synthetic subject tomograms. The difference was less than 0.0869 for 95% of all pixels and less than 0.1431 for 99% of all pixels. Since the stroma-Descemet's membrane boundary is difficult for a trained UHROCT operator to distinguish, having an average scattering potential value of 0.0869, the value of 0.0178 for 50% of all pixels is insignificant. Similarly, the endothelium-Bowman's layer boundary has a scattering potential similar to 0.1431, so a value less than 0.1431 for 99% of all pixels is visible to a trained technician. Overall, the scattering potential model fits well to the synthetic data, which was generated using the same model. The minor differences are primarily due to an imperfect structural model and to numerical issues that occur during the calculation of C.

Table 7.3: The reconstruction results show that the mean distance between the 2D corneal tomogram layer boundaries and the reconstructed corneal structure layer boundaries for the epithelium through stroma boundaries is less than 0.058 μm when using a densely spaced sagittal and transverse sampling scheme of 40 tomograms. The mean distance is less than 0.166 μm for an evenly spaced sagittal and transverse sampling scheme and less than 0.135 μm for an evenly spaced sagittal sampling scheme.

	outer-epithelium		
Synthetic Subject	dense sag. trans.	even sag. trans.	even sag.
1	0.033	0.143	0.129
2	0.094	0.210	0.124
3	0.042	0.120	0.102
4	0.035	0.136	0.129
5	0.085	0.220	0.191
mean	**0.058**	0.166	0.135
std. dev.	**0.029**	0.046	0.033

	epithelium-Bowman's layer		
Synthetic Subject	dense sag. trans.	even sag. trans.	even sag.
1	0.033	0.140	0.127
2	0.092	0.202	0.120
3	0.041	0.116	0.099
4	0.034	0.132	0.125
5	0.082	0.212	0.184
mean	**0.056**	0.160	0.131
std. dev.	**0.028**	0.044	0.032

	Bowman's layer-stroma		
Synthetic Subject	dense sag. trans.	even sag. trans.	even sag.
1	0.032	0.137	0.124
2	0.089	0.195	0.116
3	0.040	0.113	0.096
4	0.033	0.128	0.121
5	0.079	0.204	0.177
mean	**0.055**	0.155	0.127
std. dev.	**0.027**	0.041	0.030

Table 7.4: Similar to Table 7.3, the distance in μm between the 2D tomogram boundaries and the reconstructed boundaries for the stroma, Descemet's membrane, and endothelium are presented. The densely spaced sagittal and transverse sampling scheme and evenly spaced sagittal sampling scheme produce similar values for stroma to Descemet's membrane and endothelium to outer corneal boundaries.

	stroma-Descemet's membrane		
Synthetic Subject	dense sag. trans.	even sag. trans.	even sag.
1	0.028	0.049	0.039
2	0.076	0.102	0.026
3	0.010	0.020	0.013
4	0.018	0.037	0.019
5	0.048	0.093	0.072
mean	**0.036**	0.060	**0.034**
std. dev.	**0.026**	0.036	**0.023**

	endothelium-outer		
Synthetic Subject	dense sag. trans.	even sag. trans.	even sag.
1	0.028	0.048	0.037
2	0.077	0.107	0.029
3	0.011	0.023	0.015
4	0.018	0.039	0.021
5	0.050	0.099	0.077
mean	**0.037**	0.0632	**0.036**
std. dev.	**0.027**	0.038	**0.024**

Table 7.5: The scattering potential difference for the 50^{th}, 95^{th}, and 99^{th} percentile over all synthetic subjects is presented. 0.0124 is insignificant, 0.0813 is approximately the subtle change in pixel intensity used to detect the stroma-Descemet's membrane boundary, and 0.1371 is significant since it corresponds to the change in intensity for the more prominent epithelium-Bowman's membrane boundary.

	Scattering Potential Difference for all Synthetic Subjects			
Percentile	dense sag. trans.	even sag. trans.	even sag.	mean
50^{th}	0.0106	0.0178	0.0086	0.0124
95^{th}	0.0790	0.0869	0.0781	0.0813
99^{th}	0.1343	0.1431	0.1340	0.1371

7.3.2 Human cornea

Notice that while the structure appears to fit well, Fig. 7.6, any deviations from the ideal structure are apparent as abrupt edges in the scattering potential field, Fig. 7.7; this alignment is corrected through tomogram pose estimation, in Chapter 8.

Table 7.6 presents the structural reconstruction error for human corneas. On average, the reconstructed structure is within 0.18 μm of the corneal layer boundaries located on each individual tomogram. Since the reconstructed surface was fitted to the automated boundary layer locations, a low error is expected.

The difference between measured and reconstructed scattering potential, scaled between 0 and 1, is presented in Table 7.7. While these values are significantly larger than those of the synthetic cornea, which were created using the scattering potential model, the scattering potential error for 50% of all pixels is negligible, and the scattering potential for 95% of the pixels on all surfaces is less than 0.197 on average, which represents the appearance of the epithelium-Bowman's layer boundary in an UHROCT; 95% of pixels from reconstructed human subjects will have visually noticeable changes in pixel intensity. These values represent a reasonable UHROCT scattering appearance, as was designed in the scattering potential model. The model captures the measurement variations and the mean scattering potential.

(a) Subject 2 - Transverse

(e) Subject 3 - Transverse

Figure 7.6: The reconstructed corneal structure for (a) subject 2 - transverse and (b) subject 3 - transverse illustrates how corneal layer surfaces can be estimated in unsampled regions. The surface was constructed robustly despite using desired values of Θ instead of values derived from pose estimation.

134

(a) Subject 1 - Transverse

(b) Subject 3 - Sagittal

Figure 7.7: The reconstructed corneal scattering potential after removing structure for (a) subject 1 - transverse and (b) subject 2 - sagittal illustrates estimation in unsampled regions.

Table 7.6: For human subjects, the distance, in μm between the 2D tomogram boundaries and the reconstructed boundaries for the epithelium (epi.), Bowman's layer (Bow.), stroma (str), Descemet's membrane (Desc), and endothelium (end.) are presented. Like the synthetic corneal reconstruction, the difference between each layer boundary of the same subject are similar.

Subject	Sampling Dir.	outer-epi.	epi.-Bow.	Bow.-str.	str-Desc.	end-outer
1	Sagittal	0.666	0.665	0.665	0.650	0.650
1	Transversal	0.059	0.058	0.057	0.048	0.049
2	Sagittal	0.267	0.267	0.267	0.267	0.267
2	Transversal	0.032	0.032	0.032	0.030	0.030
3	Sagittal	0.065	0.064	0.063	0.056	0.056
3	Transversal	0.062	0.061	0.061	0.052	0.052
mean	Both	0.1918	0.1912	0.1908	0.1838	0.1840
std. dev.	Both	0.2476	0.2476	0.2478	0.2450	0.2449

Table 7.7: The scattering potential difference for the 50^{th}, 95^{th}, and 99^{th} percentile over all human subjects is presented. 0.044 is insignificant, but 0.197 corresponds to the change in intensity for epithelium-Bowman's membrane boundary.

	Scattering Potential Difference for Human Subjects		
Human Subject	Direction	50^{th} 95^{th}	99^{th}
1	Sagittal	0.033 0.181	0.295
1	Transversal	0.047 0.214	0.334
2	Sagittal	0.054 0.222	0.366
2	Transversal	0.035 0.169	0.278
3	Sagittal	0.065 0.218	0.359
3	Transversal	0.033 0.178	0.288
mean	Both	0.044 0.197	0.320
std. dev.	Both	0.013 0.023	0.038

Chapter 8

Corneal Tomogram Pose Estimation

The corneal eye motion is a major source of measurement location uncertainty causing the UHROCT imaging system to sample at an unintended location. This chapter attempts to estimate the tomogram sampling location using residuals, differences between the corneal measurements and measurements obtained from corneal reconstruction, satisfying Objective 6, page 45. Subsection 2.2.3 describes the processes governing eye motion and Section 6.3 derives a motion model that is essential for this chapter. Chapter 7 provides the reconstruction method that is used extensively while improving the forward model parameters. The reader should also be familiar with the visual appearance of corneal boundary layers within an UHROCT presented in Subsection 2.2.2.

Tomogram pose estimation is implemented using a large-scale optimization problem, involving six pose parameters for each tomogram and a random-field, Z_2, of size $150 \times 150 \times 150$ states, designed to translate and rotate each tomogram. Since the reconstruction process can be computationally expensive, especially for large amounts of data, the pose estimation process is designed to be iterative. At the end of each iteration, poses are selected such that the reconstruction error does not get any worse. The advantage of the iterative approach is that the optimization process, which may take several weeks to complete, can be stopped at anytime, allowing convergence to be studied during the design of the implementation.

Chapter organization

This chapter is divided into three sections. The criteria for the optimization formulation is presented in Section 8.1. Section 8.2 applies the criteria to study how well the optimization

implementation converges. The last section, Section 8.3, provides results of how the pose estimate methods performed compared to reconstruction without pose estimation.

8.1 Pose estimation criteria

The criteria that determines how well the alignment of tomograms fit to the corneal structure model and how well the measurement error from the corneal reconstruction is minimized are defined in this section. The overall optimization formulation from Section 4.3 seeks to find

$$\hat{\boldsymbol{\Theta}} = \arg_{\boldsymbol{\Theta}'} \min \varpi \big(\wp_1\left(\mathbf{m}\right) - C_1(\boldsymbol{\Theta}')\mathbf{z}_1(\boldsymbol{\Theta}')\big)^2 + (1 - \varpi)\big(\wp_2\left(\mathbf{m}\right) - C_2(\boldsymbol{\Theta}')\mathbf{z}_2(\boldsymbol{\Theta}')\big)^2.$$

(4.31 on page 61)

By incorporating mean squared structural distance, $d_{\mathrm{surf},T_i}{}^2$, and mean squared scattering measurement error, $d_{\mathrm{meas},T_i}{}^2$, between the reconstruction and the tomogram measurements, (4.31) becomes

$$\hat{\boldsymbol{\Theta}} = \arg_{\boldsymbol{\Theta}'} \min \varpi \sum_{T_i} d_{\mathrm{surf},T_i}{}^2 + (1 - \varpi) \sum_{T_i} d_{\mathrm{meas},T_i}{}^2. \tag{8.1}$$

The following subsections define the structural and scattering potential criteria.

8.1.1 Structure criterion

The structure for each tomogram, as determined using the automated layer boundary localization tool, is compared against the structure of the entire reconstruction, using the corneal structure model in Section 6.1. The objective function, (8.1) minimizes the mean squared distance between each corneal layer within each tomogram and the corresponding distance to the reconstructed layer.

Given that the reconstructed corneal structure is modelled as a spherical surface, S_ι with a radius, ρ_ι, and a center, c_ι, and the corresponding layer boundaries are represented as the curves, $\hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s)\ \forall \iota$, for tomogram T_i, the distance, $d_{\mathrm{surf},\iota,T_i}$ between a point on the spherical surface, $S_{\iota,T_i}(s)$, and the corresponding point $\hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s)$ on tomogram T_i, can be calculated using the following equations. Each point on $\hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s)$ can be transformed into a point in $\Re^3$ so that it can be compared against surface S_ι.

138

First the point $\mathbf{s}_i$ is determined from the curve $\hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s)$ for each value of s,

$$\mathbf{s}_i = \hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s). \tag{8.2}$$

Then $\mathbf{s}_i$ is transformed from the tomogram coordinate system, $\Re^2$, into the corneal system, $\Re^3$, using (4.18),

$$\mathbf{x}_i = \left(s_{0,i} - \frac{\eta_{s,0,i}}{2}\right)\frac{\gamma_{s,0,i}}{\eta_{s,0,i}}\overrightarrow{\mathbf{s}}_{0,i} + \left(s_{1,i} - \frac{\eta_{s,1,i}}{2}\right)\frac{\gamma_{s,1,i}}{\eta_{s,1,i}}\overrightarrow{\mathbf{s}}_{1,i} + \mathbf{c}_i. \tag{4.18 on page 56}$$

Using $\mathbf{x}_i$, the point on the corresponding spherical surface layer, $S_{\iota,T_i}(s)$, can be calculated as,

$$S_{\iota,T_i,0}(s) = x_{0,i} \tag{8.3}$$
$$S_{\iota,T_i,1}(s) = x_{1,i} \tag{8.4}$$
$$S_{\iota,T_i,2}(s) = \sqrt{\rho_\iota^2 - \left(S_{\iota,T_i,0}(s) - c_{\iota,0}\right)^2 - \left(S_{\iota,T_i,1}(s) - c_{\iota,1}\right)^2} + c_{\iota,2}. \tag{8.5}$$

Then $d_{\mathrm{surf},\iota,T_i}$ can be calculated for all points along $\hat{\Omega}_{\hat{\alpha}_\iota,T_i}(s)$,

$$d_{\mathrm{surf},\iota,T_i}{}^2 = \frac{\sum_s\left(S_{\iota,T_i,2}(s) - x_{2,i}\right)^2}{\sum_s 1}. \tag{8.6}$$

Then the mean distance, d_{surf,T_i} for all layers on T_i to the corresponding layers on S is

$$d_{\mathrm{surf},T_i}{}^2 = \frac{\sum_\iota\sum_s\left(S_{\iota,T_i,2}(s) - x_{2,i}\right)^2}{\sum_l\sum_s 1}. \tag{8.7}$$

Note that (8.7) weighs all layers equally by normalizing for s and that the values of $\mathbf{x}_i$ are implicity dependent on tomogram pose, $\boldsymbol{\Theta}_i$. The structural criterion is to minimize d_{surf,T_i} for tomogram T_i.

8.1.2 Scattering potential criterion

The scattering potential criterion is defined by calculating the difference in pixel intensity between each pixel in tomogram i and the corresponding pixel in the reconstructed cornea. From Section 4.1, each pixel in T_i is mapped to the measurement vector $\mathbf{m}_i$; pixel k is mapped to $m_{i,k}$. Then the corresponding position of $m_{i,k}$ in $\Re^3$ is $\mathbf{x}_{i,k}$, previously defined as

$$\mathbf{x}_{i,k} = \left(s_{0,i,k} - \frac{\eta_{s,0,i}}{2}\right)\frac{\gamma_{s,0,i}}{\eta_{s,0,i}}\overrightarrow{\mathbf{s}}_{0,i} + \left(s_{1,i,k} - \frac{\eta_{s,1,i}}{2}\right)\frac{\gamma_{s,1,i}}{\eta_{s,1,i}}\overrightarrow{\mathbf{s}}_{1,i} + \mathbf{c}_i. \tag{4.19 on page 56}$$

Using the previously defined correspondence matrix, $C_{k,j}^i$,

$$C_{k,j}^i = \begin{cases} 1 & \text{if all}\big(\mid \mathbf{x}_{i,k} - \mathbf{x}_j \mid < \gamma_{x,j}./\eta_{x,j}\big) \\ 0 & \text{else} \end{cases}, \qquad \text{(4.23 on page 58)}$$

the measurement error, $\tilde{m}_{i,k}$, can be determined from

$$\tilde{m}_{i,k} = m_{i,k} - C_k^i \mathbf{z}, \qquad (8.8)$$

where

$$C_k^i = \begin{bmatrix} C_{k,1}^i & C_{k,2}^i & \cdots & C_{k,j}^i & \cdots \end{bmatrix} \qquad (8.9)$$

The mean squared measurement error, d_{meas,T_i} for tomogram i is defined as

$$d_{\mathrm{meas},T_i}{}^2 = \frac{\sum_k \big(m_{i,k} - C_k^i \mathbf{z}\big)^2}{\sum_k 1}. \qquad (8.10)$$

Also note that C_k^i and $\mathbf{z}$ are dependent on $\boldsymbol{\Theta}_i$ and $\boldsymbol{\Theta}$ respectively.

8.2   Convergence

Using the previously defined criteria, tomogram pose estimation can now be performed using existing optimization techniques. Unfortunately, there are many nuances that determine how well the optimization converges to the desired pose estimates. To allow the optimization to converge in a reasonable amount of time, this section provides discussion and details on preliminary studies used to select an optimization method and appropriate modifications to (8.1).

A preliminary study, using data from Section 7.1, was conducted to investigate the basin of convergence [44] for individual tomogram pose parameters; the goal of this preliminary study was to determine if convergence is feasible through manipulation of a sub-set of pose parameters. For this preliminary study, each tomogram was perturbed by a given translation, $\Delta \mathbf{x}_i$, where $\Delta \mathbf{x}_i$ is proportional to the distance from the center of the sampled region. For example, if each tomogram is sampled every 20 μm in the lateral direction and each tomogram was translated by 110% of the ideal inter-spacing, then the tomogram at the center of the region would remain at its initial location, the tomogram 20 μm away from center would be translated by 2 μm, and the tomogram 40 μm away would be translated

by 4 μm. Fig. 8.1 shows the basin of convergence when all tomograms are simultaneously translated laterally as a function of translation percentage. This figure shows that, when all tomograms are translated between 80% to 120% of the ideal spacing, a gradient exists and can be calculated numerically allowing an optimization minimization function to find the ideal spacing of 100%. Since this study did not evaluate all six pose parameters for each tomogram, this example serves only as a proof of concept that convergence may be possible.

While it was desired to recalculate $\mathbf{z}$ for updated $\boldsymbol{\Theta}$ on each iteration, initial attempts at convergence for a single subject failed after running several optimization algorithms, each continuously for a week: trust-region reflective [88], trust-region dog-leg [98], and Nelder-Mead simplex search [66]. It should also be noted that Nelder-Mead marginally had the lowest residual after a week than the other methods. Investigation showed that calculating derivatives was computationally expensive. After a week, each algorithm did improve upon the initial reconstruction error, however two additional steps helped to improve convergence time.

First, an assumption was asserted that the change of $\mathbf{z}$ is negligible as long as the step size for $\boldsymbol{\Theta}$ is sufficiently small, as determined by the error associated with a quadratic model approximation [88]. Then, when the change of $\boldsymbol{\Theta}$ exceeds that step size, $\mathbf{z}$ is reconstructed using the updated $\boldsymbol{\Theta}$. However, for example, instead of creating a quadratic model [88] for $C(\boldsymbol{\Theta}')\mathbf{z}(\boldsymbol{\Theta}')$, requiring hundreds of derivatives to determine its curvature, the original function $C(\boldsymbol{\Theta}')\mathbf{z}(\boldsymbol{\Theta}')$ simply becomes $C(\boldsymbol{\Theta}')\mathbf{z}_{\boldsymbol{\Theta}''}$, where $\mathbf{z}_{\boldsymbol{\Theta}''}$ is a 0^{th} order Taylor-series approximation of $\mathbf{z}$ about $\boldsymbol{\Theta}''$. The value of $\boldsymbol{\Theta}''$ is initially selected and updated when the difference between $\boldsymbol{\Theta}'$ and $\boldsymbol{\Theta}''$ is sufficiently large. However, calculating the validity of the $C(\boldsymbol{\Theta}')\mathbf{z}_{\boldsymbol{\Theta}''}$ to specify a sufficiently large difference, is computationally expensive. For the implementation, the model, $C(\boldsymbol{\Theta}')\mathbf{z}_{\boldsymbol{\Theta}''}$, was used for all possible values of $\boldsymbol{\Theta}'$ and $\mathbf{z}_{\boldsymbol{\Theta}''}$ each time the implementation converged. Using this approach, $\mathbf{z}_{\boldsymbol{\Theta}''}$ was updated two or three times when estimating the tomogram pose for each subject.

Second, homotopy [54] was introduced to improve convergence. Instead of trying to minimize the error for all tomograms at once, the criteria for each tomogram was iteratively added during the optimization process. For instance, the method starts using only the criteria associated with tomogram 1. While the criteria for tomogram 1 is being minimized, the criteria for tomogram 2 is slowly introduced. Each additional group of criteria is introduced as the process continues, until all of the criteria has been minimized. The criteria, d^2, is calculated for each iteration, t, using the weighting functions, ϖ_{t,T_i},

$$d^2 = \sum_t \varpi_{t,T_i} \left(\varpi \sum_{T_i} d_{\text{surf},T_i}{}^2 + (1 - \varpi) \sum_{T_i} d_{\text{meas},T_i}{}^2 \right). \tag{8.11}$$
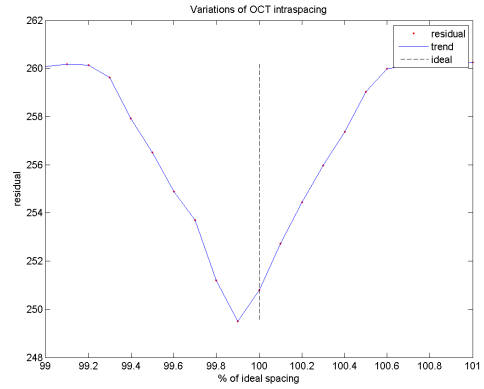
Figure 8.1: The basin of convergence when the inter-spacing between corneal layers is perturbed is illustrated for (a) synthetic subject 1 set 1, (b) synthetic subject 1 set 1 larger scale, (c) synthetic subject 2 set 1, and (d) synthetic subject 2 set 2. In all cases a gradient for the corneal inter-spacing distance exists with a minimum value near the true value of 100%.

The weighting functions ϖ_{t,T_i} are defined such that,

$$\varpi_{t,T_i} \geq \varpi_{t,T_{i+1}} \tag{8.12}$$

$$\varpi_{t+1,T_i} \geq \varpi_{t,T_i}, \tag{8.13}$$

for all t and T_i. The rate of increase for ϖ_{t,T_i} and the value of ϖ was determined experimentally until the convergence time was reduced from indefinite to several hours.

Given a synthetic cornea subject with intentionally perturbed Θ, Fig. 8.2 shows how the perturbed states converge toward the ideal states using the techniques above. While the results appear to have reached a local minima near the ideal solution, the state estimates did improve, reducing the overall reconstruction error. From this example, it should be observed that without an absolute reference point, the states tend to cluster relatively. Even though the deviation from the ideal Θ should be zero, Fig. 8.2 shows clustering about -0.05, the mean of the initial conditions, indicating that all of the tomograms are translated together from an absolute origin and that all of the tomograms are aligned with an orientation that is no longer perfectly sagittal. The resulting tomogram poses are aligned relative to each other.

Using these convergence techniques, the optimization using Nelder-Mead simplex search [66] was completed in several hours for each test subject. Algorithm 5, page 173 presents a simplified implementation of pose estimation. While the subjects are successfully reconstruction, the efficiency of the optimization implementation can be further improved in future work.

8.3   Results

Using the criteria and convergence methods defined previously in this chapter, pose estimation is applied to the human subjects specified in Section 7.1. By incorporating corneal anatomy and compensating for eye-motion, the results show that reconstruction error is successfully reduced. In addition, by observing the change of depth of each tomogram after pose estimation is applied, it is apparent that the reconstruction is consistent with corneal anatomy constraints unlike existing corneal reconstruction methods, which only align tomograms based on pixel intensity.

For human subjects, pose estimation resulted in a 32.2% to 93.6% decrease in reconstruction error from the non-pose estimated reconstruction from Chapter 7. The results for each human subject are presented in Table 8.1. In many cases, pose estimation created a better fit by accounting for motion during the scan. By comparing human subjects to the

Figure 8.2: A synthetic corneal model was perturbed such that the measured states (red) did not match the desired states. The deviation from the ideal states should be zero. After optimization, the deviation of states converged towards their final values (blue). Although they did not converge to zero, it is suspected that they converged relatively to each other since there is no absolute reference frame forcing convergence to zero. Instead, the states have clustered about a local minima, near the mean of the initial states. The reconstruction error was reduced during the minimization process.

Table 8.1: Reconstruction error for human cornea are presented for reconstruction without and with pose estimation.

| Human Subject | Direction | Mean Reconstruction Error ($\mid d^2 \mid$) | | % Error Reduction |
		Without Pose Estimation	With Pose Estimation	
1	Sagittal	0.5977	0.0425	93.6%
1	Transversal	0.0533	0.0249	58.0%
2	Sagittal	0.2447	0.0465	82.5%
2	Transversal	0.0326	0.0171	55.0%
3	Sagittal	0.0589	0.0414	32.2%
3	Transversal	0.0561	0.0249	60.3%

ideal synthetic subjects, reconstructed without any eye-motion in Section 7.3, the human subjects have the same magnitude of reconstruction error. The mean human reconstruction error, $\mid d^2 \mid$, after pose-estimation is 0.033 and the mean synthetic reconstruction error from known tomogram poses is 0.080. There is no significant difference between the human and synthetic error. This indicates that the pose estimation method is, at least partially, correcting for eye-motion and corneal anatomy.

The depth position of each tomogram for three human subject sets, illustrated in Fig. 8.3, was adjusted during pose estimation to better satisfy corneal anatomy constraints. All of the human tomogram sets were initially aligned such that their scanning depth was level. After pose estimation, the corneal curvature became consistent with expected corneal curvature. For instance, in Fig. 8.3 (a), the expected change in height, from the center of the cornea to a tomogram approximately 200 μm away, should be approximately 2.67 μm, assuming a corneal radius of approximately 7.5 mm [61], and is measured to be approximately 2.75 μm. The other subject sets, in Fig. 8.3 (b) and Fig. 8.3 (c), show similar results for their distances. Although ideal corneal anatomy should resemble an ellipse, the points associated with the depth changes are noisy. It is suspected that the pose estimation is correcting for sources of motion that contribute to the noise.

When the proposed motion corrected reconstruction results, Fig. 8.3 (a), are compared to existing tomogram stacking reconstruction methods, Fig. 8.4 (b), the difference in reconstructed corneal anatomy is significant. Without a model of anatomy, existing stacking reconstruction methods align corneal tomograms based solely on scattering potential intensity. Instead, the proposed motion corrected reconstruction integrates both corneal anatomy and scattering potential. Unlike existing methods, the proposed method

145

produces reconstructions typical of healthy corneal layer curvature.

A limitation of the optimization methodology is that the tomogram pose estimation does not appear to automatically correct for UHROCT refraction, which occurs when the imaging beam enters the cornea. Given a corneal radius of 7.5 mm [61], a distance of about 200 μm from the furthest tomogram to the center of the cornea, the angle of incidence, ϑ_{ind} is equal to

$$\vartheta_{ind} = \arcsin \frac{0.200}{7.5} = 1.53 \text{ deg.} \tag{8.14}$$

Given that the cornea has an average refractive index of 1.377 [142], from Snell's law [125], the angle of refraction, ϑ_{ref}, is equal to

$$\vartheta_{ref} = \arcsin \frac{1.377 \cdot 0.200}{7.5} = 2.10 \text{ deg.} \tag{8.15}$$

From the derivation, it is expected that the pose parameters associated with the angle of refraction should have a magnitude of about 2.10 degrees at a distance of 200 μm away from the center of the cornea. However, the measured magnitude was about 0.2 degrees, indicating that the tomogram has not been rotated enough to match the expected refraction values. It is suspected that the optimization algorithm consistently converges to a local minima instead of finding the global minimum. For example, the error caused by not correcting for refraction may be partially corrected through translation. It is also possible that the objective function does not have a minimum when the tomogram pose is consistent with the expected refraction pose. Future work should investigate these possibilities.

(a)

(b)

(c)

Figure 8.3: The change of scanning depth is corrected during pose estimation for (a) subject 3 sagittal, (b) subject 2 sagittal, and (c) subject 3 transverse. Since the initial pose is initialized to a constant depth for all transverse tomograms, a change of depth to reflect the curvature of the cornea is expected, and illustrated above, during pose estimation.

147

(a) Proposed method, reproduced from Fig. 8.3 (a)

(b) Existing stacking method

Figure 8.4: The corneal anatomy is compared between (a) the proposed reconstruction method and (b) the existing state-of-the-art tomogram stacking method. The corneal epithelium curvature of (a) the proposed method is ellipsoidal and has curvature typical of a healthy cornea. The curvature estimated using (b) existing tomogram stacking methods does not resemble an ellipsoid and has almost three times the amplitude of curvature than (a). Unlike the proposed method, existing reconstruction methods do not model or constraint corneal anatomy.

Chapter 9

Conclusions

This chapter describes the significance of chapters 5 through 8. The impact of corneal reconstruction and corneal layer thickness measurements are presented in addition to other projects that have utilized statistical modelling based on the techniques in this document.

9.1 Objectives

This subsection concludes on the status of each objective introduced in Chapter 1 and detailed in Chapter 3. The success of each objective shows that the proposed reconstruction methodology is the first method to not only combine corneal anatomy and UHROCT imaging into a 3D model, but to also use the model to successfully reconstruct corneal structure and to compensate for eye-motion. Unlike existing reconstruction methods, the proposed method generates reconstructions that are consistent with healthy human anatomy.

1. **Extend general reconstruction theory**
   This objective focuses on the development of a systematic method capable of producing a 3D reconstruction from a set of tomograms. The method accounts for measurement error and unknown tomogram pose. Chapter 4 develops the framework and supporting mathematics utilized throughout the remaining chapters to fulfill this objective. The reconstruction results in a corneal application area, objectives 5 and 6, provide validation of the proposed method.

2. **Create 2D corneal structure model**
   A structural model capable of parameterizing a 2D corneal tomogram is derived in

Chapter 5. This model was validated using twelve human subjects, thousands of tomograms, through the implementation of a method that can automatically locate each corneal layer boundary to within 6.52 μm. This method was the first method to locate five corneal layer boundaries from UHROCT. The 2D model allowed layers to be located despite prominent image artefacts that thwarted other non-model based methods.

3. **Create 3D corneal model**

   Chapter 6 extends the 2D corneal structural model into a 3D corneal model that combines both structure and scattering potential. A parametric 3D structural surface was designed based on the statistics extracted using the 2D corneal structure model. Several parametric models were proposed, including rotated polynomials and spherical models, and are validated in objectives 5 and 6. Eye rotation is also modelled and is utilized to produce synthetic tomograms from a synthetic cornea generated from this model.

4. **Generate synthetic cornea from model**

   The parameters of the 3D corneal model can be manipulated to generate a synthetic model. Each synthetic cornea contains a 3D random field of scattering potential data and a parametric surface for each corneal layer boundary. The document has shown how tomograms are sampled from the synthetic model for use in the development of reconstruction methods. The success of this objective allows reconstruction to be tested on synthetic data containing known ground truth.

5. **Reconstruct cornea from tomograms without tomogram motion correction**

   Chapter 7 applies the methodology of objective 1 to corneal reconstruction using the corneal model created in objective 3. The reconstruction method was proven using synthetic cornea with known ground truth created through the success of objective 4. The scattering potential reconstruction was proven by reconstructing tomograms sampled from the synthetic cornea and comparing the results against the known, synthesized, scattering potential. Similar results were shown for estimating synthetic corneal structure. As expected, the results of the synthetic testing showed minimal differences between the reconstructed structure and scattering potential and the corresponding synthesized values. Having been proven for synthetic corneas, the reconstruction tests were repeated for manually segmented human corneas, which also provided reasonable reconstruction error considering that the reconstruction did not explicitly account for eye-motion during the corneal imaging process. To the author's knowledge this is the first successful method to include anatomy to create a 3D corneal reconstruction from tomograms.

6. **Reconstruct cornea from tomograms with tomogram motion correction**

In addition to using a corneal model to distinguish this corneal reconstruction method from existing state of the art, the reconstruction is further extended to translate and orient each tomogram to better align the tomograms with corneal anatomy while simultaneously aligning tomograms based on scattering potential. Pose-estimation using multi-objective optimization criteria successfully illustrates how the tomogram alignment after reconstruction becomes consistent with corneal anatomy constraints; the reconstruction estimated the corneal curvature of the subject and aligned tomograms to correspond. Results on human subjects are greatly improved when compared to the non-motion corrected reconstruction in objective 5. While convergence can be improved, a means to correct tomogram pose is validated.

7. **Extract corneal layer thickness measurements**

By the design of the corneal model, corneal layer thickness measurements can be easily and directly obtained at any point within the reconstruction since the corneal structure is parameterized and decoupled from the scattering potential measurements. One advantage of using the structure model is that the position of corneal layer boundaries can be estimated from sparse sets of tomograms that require less time to obtain than dense sampling required by existing reconstruction methods. Imaging sessions are shorter and, consequently, there is less error due to eye-motion. The result of the successful reconstruction in objectives 5 and 6 is that corneal layer thickness can be measured at any point within the parameterized model.

9.2   Impact and contributions

While the previous section explicitly discusses each objective, this section describes the overall impact of this research. The reconstruction methods presented in this document utilize models of corneal anatomy, scattering potential, and eye-motion. This is a significant improvement over existing literature, since many of the existing corneal reconstruction techniques stack corneal tomograms using alignment methods based solely on the scattering potential intensity.

9.2.1   3D model-based corneal reconstruction

Even if motion is not accounted for, such as in the 3D reconstruction from Chapter 7, unmeasured corneal regions, those not directly measured through imaging, can be esti-

mated from the corneal model. In particular, the structural model allows researchers to perform corneal layer thickness measurements at unsampled locations. Unlike existing reconstruction, producing tomogram stacks from densely sampled corneal regions, the proposed reconstruction allows sparse sampling of the corneal region as long as the epithelium and endothelium layer boundaries are visible.

When motion is accounted for through tomogram pose estimation, Chapter 8, the reconstruction error is further reduced as the estimates of tomogram pose are improved. During pose estimation, the structural constraints encourage the tomograms to translate and reorient so that tomograms better follow the curvature of the cornea and, using the scattering potential model, pose estimation attempts to reduce differences between scattering potential for each tomogram and its corresponding value in the reconstruction.

9.2.2 Layer thickness measurements

In conjunction with reconstruction, this document presents the first automated corneal layer boundary localization algorithm capable of segmenting five corneal layer boundaries from OCT tomograms. The proposed segmentation algorithm can automatically generate thickness measurements from each segmented layer. Using the structural correspondence of the epithelium and endothelium and their spatial relationship with the remaining layers, the locations of the five corneal layers can be calculated. The experimental results from Section 5.5 demonstrate the robustness and applicability of the method applied to over $2,050$ images.

The automatic corneal layer boundary localization is operationally used in the Department of Optometry, University of Waterloo to locate and measure the thickness of corneal layers [10,57]. The department has used measurements from this method to study the relationship between contact lens induced hypoxia with corneal swelling. Other potential uses of this work include the segmentation of tissue layers that occur in medical and biological imaging or thickness measurements of sediment core samples taken from lakes. In particular, this framework is applicable in situations where measuring the thickness between two correlated layers is important.

9.2.3 Other statistical modelling contributions

There are other non-corneal related statistical modelling projects based on the same principals used for corneal modelling. Ongoing collaboration with the University of Waterloo Anthropology Department has helped to develop a statistical model that can be used to

determine the species and sex of olive and yellow baboons based on nasal cavity measurements obtained from computed tomography (CT) imaging of baboon skulls. The hope is that trends found in baboons can be correlated to recent human ancestors to identify hybrids between Neanderthals and anatomically modern humans.

Another project utilizes statistical modelling to reduce error associated with a local positioning system [36]. This system was designed to locate a transponder, primarily using vision. A set of accelerometers and gyroscopes were incorporated to continue tracking if the vision system was obstructed. Human movement was modelled as a hidden Markov model to incorporate measurements with the best signal processing and filtering given a specific type of human action. Based on the statistical modeling aspects for the accelerometer and other modelling aspects for the vision system, this work has won three international industrial awards, the PLASA 2012 Gold Innovation Award, IABM 2012 Game-Changer Award, and the PLASA 2010 Innovation Award. Each award was granted by committees for technical excellence and potential to change an industry.

9.3 Future work

The ability to produce large, three-dimensional corneal reconstructions facilitates significant clinical and research opportunities. Collaborators in science and optometry are eager to continue refining the methodology so that fine structural details of corneal and retinal anatomy can be revealed and modelled. One of the major issues is the efficiency of the current implementation. While much of the implementation tried to take advantage of Matlab best practices, e.g. avoiding for-loops, creating custom classes to represent matrices as non-copied handles, and avoiding computationally expensive visualizations, these practices have not been applied in all circumstances, leaving room for additional improvements. For instance, simply loading a large human OCT can take up to ten minutes using the proprietary Matlab format; some of the memory paging techniques used to load independent texture data into memory during texture synthesis might be applied to save and load OCTs. If this particular corneal reconstruction implementation is to be used in future studies, these computational limitations should be improved.

Further, some tasks are highly parallelizable and can be better solved using hundreds and thousands of low-powered processors, such as those found in graphics cards. Since calculating the position correspondence between $\mathbf{m}$ and $\mathbf{z}$ is highly parallelizable, but computationally expensive for large volumes, calculating the correspondence matrix C on the parallel computing framework is ideal. While some aspects, such as texture generation iterations, already utilize a graphics card, many such improvements can be made.

The resolution of the reconstructed area can also be increased as the efficiency improves. One ongoing project is modelling the development of shark embryo growth. The study analyzes high resolution 2D images of cellular development over time. Unfortunately the reduction of resolution during the spatial and temporal reconstruction prevents the fine cellular boundaries of the shark embryo from being located. The best approach is to increase efficiency and to introduce additional reconstruction decoupling so that different spatial and temporal regions can be reconstructed first and then stitched together to form a larger, consistent reconstruction. The goal is to subdivide the reconstruction area into segments that do not need to exist simultaneously in memory. In a similar sense, corneal patient data can be reconstructed over time as well. The key is to have sufficient resolution to capture important measurements for the application.

9.4 Summary

Overall, the corneal model and reconstruction successfully reconstruct human and synthetic corneal tomograms. The automated corneal layer boundary localization was proven after successfully segmenting over $2,000$ tomograms and proven in studies that measure change in corneal layer thickness. With improvements to computational efficiency and increased reconstruction resolution, the methodology can be applied to model finely detailed cellular structures in future applications such as studying the development of shark embryo cells and the effects of pathology on corneal layer cells.

References

[1] R. R. Ackermann, J. Rogers, and J. M. Cheverud. Identifying the morphological signatures of hybridization in primate and human evolution. *Journal of Human Evolution*, 51(6):632–645, 2006.

[2] O. Ahmad, K. Ramamurthi, K. Wilson, K. Engelke, R. Prince, and R. Taylor. Volumetric dxa (vxa): A new method to extract 3d information from multiple in vivo dxa images. *Journal of Bone and Mineral Research*, 25(12):2744–2751, 2010.

[3] A. A. Amini, T. E. Weymouth, and R. C. Jain. Using dynamic programming for solving variational problems in vision. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 12(9):855–867, 1990.

[4] G. Baikoff, E. Lutun, C. Ferraz, and J. Wei. Static and dynamic analysis of the anterior segment with optical coherence tomography. *Journal of Cataract & Refractive Surgery*, 30(9):1843–1850, 2004.

[5] B. A. Barsky and A. D. DeRose. Geometric continuity of parametric curves. Technical report, University of California at Berkeley, Berkeley, CA, USA, 1984.

[6] J. Bartlett and S. Jaanus. *Clinical Ocular Pharmacology*. Butterworth-Heinemann / Elsevier, 2008.

[7] W. Bates. *The Cure of Imperfect Sight by Treatment Without Glasses*. CreateSpace, 2011.

[8] M. Bechmann, M. J. Thiel, A. S. Neubauer, S. Ullrich, K. Ludwig, K. R. Kenyon, and M. W. Ulbig. Central corneal thickness measurement with a retinal optical coherence tomography device versus standard ultrasonic pachymetry. *Cornea*, 20(1):50–54, 2001.

[9] O. Bergamin, S. Bizzarri, and D. Straumann. Ocular torsion during voluntary blinks in humans. *Investigative Ophthalmology & Visual Science*, 43(11):3438–3443, 2002.

[10] K. Bizheva, C. Hyun, J. A. Eichel, S. Hariri, A. Mishra, D. A. Clausi, P. Fieguth, T. Simpson, and N. Hutchings. Evaluation of hypoxic swelling of human cornea with high speed ultrahigh resolution optical coherence tomography. In *SPIE BiOS: Biomedical Optics*, pages 71631G–71631G. International Society for Optics and Photonics, 2009.

[11] O. N. Bjornstad, W. Falck, and N. C. Stenseth. A geographic gradient in small rodent density fluctuations: a statistical modelling approach. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 262(1364):127–133, 1995.

[12] C. E. Buck, W. G. Cavanagh, and C. D. Litton. *Bayesian Approach to Interpreting Archaeological Data*. Wiley Chichester, 1996.

[13] R. Butt. *Introduction to Numerical Analysis Using MATLAB?* Infinity Science Series. Jones & Bartlett Learning, 2009.

[14] L. Bye, N. Modi, and M. Stanford. *Basic Sciences for Ophthalmology*. Oxford Specialty Training: Basic Science. OUP Oxford, 2013.

[15] J. Cameron, J. Skofronick, and R. Grant. *Physics of the Body*. Medical Physics Series. Medical Physics Publishing Corporation, 1999.

[16] R. H. Carpenter. *Movements of the Eyes*. Pion Limited, 1988.

[17] T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Transaction on Image Processing*, 10(2):266–277, 2001.

[18] T. C. Chen, B. Cense, M. C. Pierce, N. Nassif, B. H. Park, S. H. Yun, B. R. White, B. E. Bouma, G. J. Tearney, and J. F. de Boer. Spectral domain optical coherence tomography: ultra-high speed, ultra-high resolution ophthalmic imaging. *Archives of Ophthalmology*, 123(12):1715, 2005.

[19] V. Christopoulos, L. Kagemann, G. Wollstein, H. Ishikawa, M. L. Gabriele, M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, J. S. Duker, D. K. Dhaliwal, et al. In vivo corneal high-speed, ultra high-resolution optical coherence tomography. *Archives of Ophthalmology*, 125(8):1027, 2007.

[20] R. T. Clarke et al. *Statistical Modelling in Hydrology*. John Wiley & Sons, 1994.

156

[21] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, 1991.

[22] A. C. F. Colchester and D. J. Hawkes, editors. *IPMI '91: Proceedings of the 12th International Conference on Information Processing in Medical Imaging*, London, UK, UK, 1991. Springer-Verlag.

[23] P. Congdon. *Bayesian Statistical Modelling*. Wiley. com, 2007.

[24] A. R. Conn, N. I. Gould, and P. L. Toint. *Trust Region Methods*. Number 1. Siam, 2000.

[25] C. Cook and M. Langham. Corneal thickness in interstitial keratitis. *The British Journal of Ophthalmology*, 37(5):301, 1953.

[26] R.-P. Copt, R. Thomas, and A. Mermoud. Corneal thickness in ocular hypertension, primary open-angle glaucoma, and normal tension glaucoma. *Archives of Ophthalmology*, 117(1):14, 1999.

[27] G. R. Cross and A. K. Jain. Markov random field texture models. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 5(1):25, 1983.

[28] P. Dani and S. Chaudhuri. Automated assembling of images: Image montage preparation. *Pattern Recognition*, 28(3):431–445, 1995.

[29] B. Delaunay. Sur la sphere vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk*, 7(793-800):1–2, 1934.

[30] F. C. Delori, R. H. Webb, and D. H. Sliney. Maximum permissible exposures for ocular safety (ansi 2000), with emphasis on ophthalmic devices. *Journal of the Optical Society of America A*, 24(5):1250–1265, 2007.

[31] R. Ditchburn and B. Ginsborg. Involuntary eye movements during fixation. *The Journal of Physiology*, 119(1):1–17, 1953.

[32] W. Drexler and J. G. Fujimoto. *Optical Coherence Tomography: Technology and Applications*. Springer, 2008.

[33] W. Drexler, U. Morgner, R. K. Ghanta, F. X. Kärtner, J. S. Schuman, and J. G. Fujimoto. Ultrahigh-resolution ophthalmic optical coherence tomography. *Nature Medicine*, 7(4):502–507, 2001.

[34] H. S. Dua, L. A. Faraj, D. G. Said, T. Gray, and J. Lowe. Human corneal anatomy redefined: A novel pre-descemet's layer (dua's layer). *Ophthalmology*, 2013.

[35] J. A. Eichel, K. K. Bizheva, D. A. Clausi, and P. W. Fieguth. Automated 3d reconstruction and segmentation from optical coherence tomography. In *Computer Vision–ECCV 2010*, pages 44–57. Springer, 2010.

[36] J. A. Eichel, D. A. Clausi, and P. Fieguth. Precise high speed multi-target multi-sensor local positioning system. In *Computer and Robot Vision (CRV), 2011 Canadian Conference on*, pages 109–116. IEEE, 2011.

[37] R. W. Emerson. *Selected Writings of Ralph Waldo Emerson*. Penguin. com, 2003.

[38] L. Fahrmeir, G. Tutz, and W. Hennevogl. *Multivariate Statistical Modelling based on Generalized Linear Models*, volume 2. Springer New York, 1994.

[39] J. Feng and H. Ip. A multi-resolution statistical deformable model (misto) for soft-tissue organ reconstruction. *Pattern Recognition*, 42(7):1543–1558, 2009.

[40] A. F. Fercher, W. Drexler, C. K. Hitzenberger, and T. Lasser. Optical coherence tomography-principles and applications. *Reports on Progress in Physics*, 66(2):239, 2003.

[41] A. F. Fercher et al. Optical coherence tomography. *Journal of Biomedical Optics*, 1(2):157–173, 1996.

[42] P. Fieguth. *Statistical Image Processing and Multidimensional Modeling*. Springer, 2010.

[43] G. R. Fishman, M. E. Pons, J. A. Seedor, J. M. Liebmann, and R. Ritch. Assessment of central corneal thickness using optical coherence tomography. *Journal of Cataract & Refractive Surgery*, 31(4):707–711, 2005.

[44] C. Floudas and P. Pardalos. *Encyclopedia of Optimization*. Number v. 1 in Encyclopedia of Optimization. Springer, 2008.

[45] J. Forrester, A. Dick, P. McMenamin, W. Lee, and L. Boyev. The eye: Basic sciences and practices. *American Journal of Ophthalmology*, 122(6):918–918, 1996.

[46] J. Frank. Three-dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state. *New York*, 2006.

[47] A. Frisch, B. Hnich, Z. Kiziltan, I. Miguel, and T. Walsh. Global constraints for lexicographic orderings. In *Principles and Practice of Constraint Programming-CP 2002*, pages 93–108. Springer, 2006.

[48] M. K. Garvin, M. D. Abràmoff, X. Wu, S. R. Russell, T. L. Burns, and M. Sonka. Automated 3-d intraretinal layer segmentation of macular spectral-domain optical coherence tomography images. *Medical Imaging, IEEE Transactions on*, 28(9):1436–1447, 2009.

[49] E. M. Gertz. A quasi-newton trust-region method. *Mathematical Programming*, 100(3):447–470, 2004.

[50] J. N. Goldman and G. B. Benedek. The relationship between morphology and transparency in the nonswelling corneal stroma of the shark. *Investigative Ophthalmology & Visual Science*, 6(6):574–600, 1967.

[51] M. O. Gordon, J. A. Beiser, J. D. Brandt, D. K. Heuer, E. J. Higginbotham, C. A. Johnson, J. L. Keltner, J. P. Miller, I. Parrish, K. Richard, et al. The ocular hypertension treatment study: baseline factors that predict the onset of primary open-angle glaucoma. *Archives of Ophthalmology*, 120(6):714, 2002.

[52] J. Gu and P. Bourne. *Structural Bioinformatics*. Methods of Biochemical Analysis. Wiley, 2009.

[53] J. Hajnal and D. Hill. *Medical Image Registration*. Biomedical Engineering. Taylor & Francis, 2010.

[54] J.-H. He. Homotopy perturbation technique. *Computer methods in applied mechanics and engineering*, 178(3):257–262, 1999.

[55] E. Holland, M. Mannis, and W. Lee. *Ocular Surface Disease: Cornea, Conjunctiva and Tear Film*. Elsevier Health Sciences, 2013.

[56] D. Huang, E. A. Swanson, C. P. Lin, J. S. Schuman, W. G. Stinson, W. Chang, M. R. Hee, T. Flotte, K. Gregory, C. A. Puliafito, et al. Optical coherence tomography. *Science*, 254(5035):1178–1181, 1991.

[57] N. Hutchings, T. L. Simpson, C. Hyun, A. A. Moayed, S. Hariri, L. Sorbara, and K. Bizheva. Swelling of the human cornea revealed by high-speed, ultrahigh-resolution optical coherence tomography. *Investigative Ophthalmology & Visual Science*, 51(9):4579–4584, 2010.

159

[58] R. J. Jacob. The use of eye movements in human-computer interaction techniques: what you look at is what you get. *ACM Transactions on Information Systems (TOIS)*, 9(2):152–169, 1991.

[59] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.

[60] R. Kindermann, J. L. Snell, et al. *Markov Random Fields and their Applications*. American Mathematical Society Providence, RI, 1980.

[61] S. A. Klein. Corneal topography: A review, new ansi standards and problems to solve. In *Vision Science and its Applications*. Optical Society of America, 2000.

[62] A. Konstantopoulos, G. Yadegarfar, M. Fievez, D. F. Anderson, and P. Hossain. In vivo quantification of bacterial keratitis with optical coherence tomography. *Investigative Ophthalmology & Visual Science*, 52(2):1093–1097, 2011.

[63] D. Koozekanani, K. Boyer, and C. Roberts. Retinal thickness measurements from optical coherence tomography using a markov boundary model. *IEEE Transactions on Medical Imaging*, 20(9):900–916, 2001.

[64] T. Krishnamurti, C. Kishtawal, T. E. LaRow, D. R. Bachiochi, Z. Zhang, C. E. Williford, S. Gadgil, and S. Surendran. Improved weather and seasonal climate forecasts from multimodel superensemble. *Science*, 285(5433):1548–1550, 1999.

[65] J. Kuo. *Electron Microscopy: Methods and Protocols*. Methods in molecular biology. Humana Press, 2007.

[66] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright. Convergence properties of the nelder–mead simplex method in low dimensions. *SIAM Journal on Optimization*, 9(1):112–147, 1998.

[67] F. LaRocca, S. J. Chiu, R. P. McNabb, A. N. Kuo, J. A. Izatt, and S. Farsiu. Robust automatic segmentation of corneal layer boundaries in sdoct images using graph theory and dynamic programming. *Biomedical Optics Express*, 2(6):1524, 2011.

[68] B. H. Lee, J. Liu, Z. Tan, E. Heng, J. Cheng, N. M. Tan, D. W. K. Wong, E. Trucco, J. Mehta, and T. Y. Wong. Corneal graft detection for descemet's stripping automated endothelial keratoplasty using optical coherence tomography. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*, pages 3037–3040. IEEE, 2010.

[69] A. Lens, S. C. Nemeth, and J. K. Ledford. *Ocular Anatomy and Physiology*. Slack Incorporated, 2008.

[70] L. Levin, S. Nilsson, J. Hoeve, S. Wu, P. Kaufman, and A. Alm. *Adler's Physiology of the Eye: Expert Consult*. Elsevier Health Sciences, 2011.

[71] H. F. Li, W. M. Petroll, T. Møller-Pedersen, J. K. Maurer, H. D. Cavanagh, and J. V. Jester. Epithelial and corneal thickness measurements by in vivo confocal microscopy through focusing (cmtf). *Current Eye Research*, 16(3):214–221, 1997.

[72] S. Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag New York, Inc., 1995.

[73] Y. Li, D. M. Meisler, M. Tang, A. T. Lu, V. Thakrar, B. J. Reiser, and D. Huang. Keratoconus diagnosis with optical coherence tomography pachymetry mapping. *Ophthalmology*, 115(12):2159–2166, 2008.

[74] Y. Li, R. Shekhar, and D. Huang. Corneal pachymetry mapping with high-speed optical coherence tomography. *Ophthalmology*, 113(5):792–799, 2006.

[75] G. Lukács, R. Martin, and D. Marshall. Faithful least-squares fitting of spheres, cylinders, cones and tori for reliable segmentation. In *Computer VisionECCV'98*, pages 671–686. Springer, 1998.

[76] R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995.

[77] C. M. Martín-Navarro, J. Lorenzo-Morales, M. G. Cabrera-Serra, F. Rancel, N. M. Coronado-Álvarez, J. E. Piñero, and B. Valladares. The potential pathogenicity of chlorhexidine-sensitive acanthamoeba strains isolated from contact lens cases from asymptomatic individuals in tenerife, canary islands, spain. *Journal of Medical Microbiology*, 57(11):1399–1404, 2008.

[78] S. Martinez-Conde, S. L. Macknik, and D. H. Hubel. The role of fixational eye movements in visual perception. *Nature Reviews Neuroscience*, 5(3):229–240, 2004.

[79] B. R. Masters and S. W. Paddock. Three-dimensional reconstruction of the rabbit cornea by confocal scanning optical microscopy and volume rendering. *Applied Optics*, 29(26):3816–3822, 1990.

[80] A. Mathai. *Jacobians of Matrix Transformations and Functions of Matrix Argument.* World Scientific Pub., 1997.

[81] D. M. Maurice. The structure and transparency of the cornea. *The Journal of Physiology*, 136(2):263, 1957.

[82] F. A. Medeiros, P. A. Sample, L. M. Zangwill, C. Bowd, M. Aihara, and R. N. Weinreb. Corneal thickness as a risk factor for visual field loss in patients with preperimetric glaucomatous optic neuropathy. *American Journal of Ophthalmology*, 136(5):805–813, 2003.

[83] M. Millodot. Dictionary of optometry and visual science, 7th. *Clinical and Experimental Optometry*, 92, 2009.

[84] A. Mishra, A. Wong, K. Bizheva, and D. A. Clausi. Intra-retinal layer segmentation in optical coherence tomography images. *Optics Express*, 17(26):23719–23728, 2009.

[85] A. Mishra, A. Wong, W. Zhang, D. A. Clausi, and P. Fieguth. Improved interactive medical image segmentation using enhanced intelligent scissors. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE*, pages 3083–3086. IEEE, 2008.

[86] A. A. Moayed, S. Hariri, E. S. Song, V. Choh, and K. Bizheva. In vivo volumetric imaging of chicken retina with ultrahigh-resolution spectral domain optical coherence tomography. *Biomedical Optics Express*, 2(5):1268–1274, 2011.

[87] F. Møller, M. Laursen, J. Tygesen, and A. Sjølie. Binocular quantification and characterization of microsaccades. *Graefe's Archive for Clinical and Xxperimental Ophthalmology*, 240(9):765–770, 2002.

[88] J. J. Moré and D. C. Sorensen. Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572, 1983.

[89] L. Moreaux, O. Sandre, S. Charpak, M. Blanchard-Desce, and J. Mertz. Coherent scattering in multi-harmonic light microscopy. *Biophysical Journal*, 80(3):1568–1574, 2001.

[90] E. N. Mortensen and W. A. Barrett. Intelligent scissors for image composition. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, pages 191–198. ACM, 1995.

[91] S. Muscat, N. McKay, S. Parks, E. Kemp, and D. Keating. Repeatability and re-producibility of corneal thickness measurements by optical coherence tomography. *Investigative Ophthalmology & Visual Science*, 43(6):1791–1795, 2002.

[92] J. Parker and P. Parker. *The Official Patient's Sourcebook on Fuchs' Dystrophy.* Official Patient Guides. ICON Health Publications, 2002.

[93] E. E. Pastorino and S. M. Doyle-Portillo. *What is Psychology?* CengageBrain.com, 2011.

[94] D. Pavan-Langston. *Manual of Ocular Diagnosis and Therapy.* Lippincott Manual Series. Lippincott Williams & Wilkins, 2008.

[95] M. Pavelka and J. Roth. *Functional Ultrastructure: Atlas of Tissue Biology and Pathology.* Springer, 2010.

[96] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld. Adaptive histogram equalization and its variations. *Computer Vision, Graphics, and Image Processing*, 39(3):355–368, 1987.

[97] B. Potsaid, I. Gorczynska, V. J. Srinivasan, Y. Chen, J. Jiang, A. Cable, and J. G. Fujimoto. Ultrahigh speed spectral/fourier domain oct ophthalmic imaging at 70,000 to 312,500 axial scans per second. *Optics Express*, 16(19):15149, 2008.

[98] W. H. Press. *Numerical Recipes in Fortran 77: The Art of Scientific Computing.* Cambridge university press, 1992.

[99] J. M. Prewitt. *Object Enhancement and Extraction.* Academic Press, New York, 1970.

[100] R. M. Pritchard. Stabilized images on the retina. *Scientific American*, 1961.

[101] P. Puvanathasan, P. Forbes, Z. Ren, D. Malchow, S. Boyd, and K. Bizheva. High-speed, high-resolution fourier-domain optical coherence tomography system for reti-nal imaging in the 1060 nm wavelength region. *Optics Letters*, 33(21):2479–2481, 2008.

[102] S. Radhakrishnan, J. Goldsmith, D. Huang, V. Westphal, D. K. Dueker, A. M. Rollins, J. A. Izatt, and S. D. Smith. Comparison of optical coherence tomogra-phy and ultrasound biomicroscopy for detection of narrow anterior chamber angles. *Archives of Ophthalmology*, 123(8):1053, 2005.

[103] J. L. B. Ramos, Y. Li, and D. Huang. Clinical and research applications of anterior segment optical coherence tomography–a review. *Clinical & Experimental Ophthalmology*, 37(1):81–89, 2009.

[104] R. Rardin. *Optimization in Operations Research*. Prentice-Hall International, 1998.

[105] F. Ratliff and L. A. Riggs. Involuntary motions of the eye during monocular fixation. *Journal of Experimental Psychology*, 40(6):687, 1950.

[106] T. Reinhard and D. Larkin. *Cornea and External Eye Disease: Corneal Allotransplantation, Allergic Disease and Trachoma*. Essentials in Ophthalmology. Springer, 2010.

[107] T. Reinhard and F. Larkin. *Corneal Disease: Recent Developments in Diagnosis and Therapy*. Springer, 2012.

[108] D. Z. Reinstein, T. J. Archer, M. Gobbe, R. H. Silverman, and D. J. Coleman. Stromal thickness in the normal cornea: three-dimensional display with artemis very high-frequency digital ultrasound. *Journal of Refractive Surgery*, 25(9):776, 2009.

[109] M. Riazi-Esfahani, G. A. Peyman, E. Aydin, A. A. Kazi, M. Kivilcim, and D. R. Sanders. Prevention of corneal neovascularization: evaluation of various commercially available compounds in an experimental rat model. *Cornea*, 25(7):801–805, 2006.

[110] L. A. Riggs, J. C. Armington, and F. Ratliff. Motions of the retinal image during fixation. *Journal of the Optical Society of America*, 44(4):315–321, 1954.

[111] M. Rosenfield, N. Logan, and K. Edwards. *Optometry: Science, Techniques and Clinical Management*. Butterworth Heinemann Elsevier, 2009.

[112] Y. A. Rozanov. *Markov Random Fields*. Springer, 1982.

[113] T. Sauer. *Numerical Analysis*. Springer, New York, 2011.

[114] F. Scarpa, D. Fiorin, and A. Ruggeri. In vivo three-dimensional reconstruction of the cornea from confocal microscopy images. In *Engineering in Medicine and Biology Society, 2007. EMBS 2007. 29th Annual International Conference of the IEEE*, pages 747–750. IEEE, 2007.

[115] M. Schlüter. Analysis of holographic interferograms with a tv picture system. *Optics & Laser Technology*, 12(2):93–95, 1980.

164

[116] J. Schmitt and S. Xiang. Cross-polarized backscatter in optical coherence tomography of biological tissue. *Optics Letters*, 23(13):1060–1062, 1998.

[117] T. Schmoll, A. Unterhuber, C. Kolbitsch, T. Le, A. Stingl, and R. Leitgeb. Precise thickness measurements of bowman's layer, epithelium, and tear film. *Optometry & Vision Science*, 89(5):E795–E802, 2012.

[118] R. Serway, R. Beichner, and J. Jewett. *Physics for Scientists and Engineers*. Number v. 1 in Physics for Scientists and Engineers. Saunders College Publishing, 2000.

[119] F. Shih. *Image Processing and Mathematical Morphology: Fundamentals and Applications*. Taylor & Francis, 2010.

[120] K. Shoemake. Animating rotation with quaternion curves. *ACM SIGGRAPH Computer Graphics*, 19(3):245–254, 1985.

[121] M. A. Shousha, V. L. Perez, J. Wang, T. Ide, S. Jiao, Q. Chen, V. Chang, N. Buchser, S. R. Dubovy, W. Feuer, et al. Use of ultra-high-resolution optical coherence tomography to detect in vivo characteristics of descemet's membrane in fuchs' dystrophy. *Ophthalmology*, 117(6):1220–1227, 2010.

[122] T. Simpson and D. Fonn. Optical coherence tomography of the anterior segment. *The Ocular Surface*, 6(3):117–127, 2008.

[123] A. F. Smith. A general bayesian linear model. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 67–75, 1973.

[124] G. Smolin, M. C Stephen Foster, D. T. Azar, and C. H. Dohlman. *Smolin and Thoft's the Cornea: Scientific Foundations and Clinical Practice*. Wolters Kluwer Health, 2005.

[125] R. S. Snell, M. A. Lemp, and I. Grunther. *Clinical Anatomy of the Eye*. Blackwell Science, 1998.

[126] A. H. Solberg, T. Taxt, and A. K. Jain. A markov random field model for classification of multisource satellite imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 34(1):100–113, 1996.

[127] D. C. Sorensen. Newton's method with a model trust region modification. *SIAM Journal on Numerical Analysis*, 19(2):409–426, 1982.

[128] R. Srebro. Fixation of normal and amblyopic eyes. *Archives of Ophthalmology*, 101(2):214, 1983.

[129] O. Stachs, A. Zhivov, R. Kraak, J. Stave, and R. Guthoff. In vivo three-dimensional confocal laser scanning microscopy of the epithelial nerve structure in the human cornea. *Graefe's Archive for Clinical and Experimental Ophthalmology*, 245(4):569–575, 2007.

[130] D. Stocum. *Regenerative Biology and Medicine*. Academic Press, 2012.

[131] V. Sundarapandian. *Numerical Linear Algebra*. PHI Learning, 2008.

[132] L. Szczotka-Flynn, D. G. Ahearn, J. Barr, W. J. Benjamin, T. Kiang, J. J. Nichols, O. D. Schein, R. P. Stone, and L. Winterton. History, evolution, and evolving standards of contact lens care. *Contact Lens and Anterior Eye*, 36:S4–S8, 2013.

[133] E. Traboulsi. *Genetic Diseases of the Eye*. Oxford Monographs on Medical Genetics. Oxford University Press, USA, 2011.

[134] L. Vabre, A. Dubois, and A. C. Boccara. Thermal-light full-field optical coherence tomography. *Optics Letters*, 27(7):530–532, 2002.

[135] G. Van de Wouwer, P. Scheunders, and D. Van Dyck. Statistical texture characterization from discrete wavelet representations. *IEEE Transaction on Image Processing*, 8(4):592–598, 1999.

[136] R. Verma. *Wave Optics*. Discovery Publishing House Pvt. Limited, 2006.

[137] E. N. Vithana, T. Aung, C. C. Khor, B. K. Cornes, W.-T. Tay, X. Sim, R. Lavanya, R. Wu, Y. Zheng, M. L. Hibberd, et al. Collagen-related genes influence the glaucoma risk factor, central corneal thickness. *Human Molecular Genetics*, 20(4):649–658, 2011.

[138] J. Wang, J. Thomas, I. Cox, and A. Rollins. Noncontact measurements of central corneal epithelial and flap thickness after laser in situ keratomileusis. *Investigative Ophthalmology & Visual Science*, 45(6):1812–1816, 2004.

[139] L. Wilkins. *Professional Guide to Diseases*. LWW medical book collection. Wolters Kluwer Health/Lippincott Williams & Wilkins, 2009.

[140] M. Wojtkowski, V. Srinivasan, J. G. Fujimoto, T. Ko, J. S. Schuman, A. Kowalczyk, and J. S. Duker. Three-dimensional retinal imaging with high-speed ultrahigh-resolution optical coherence tomography. *Ophthalmology*, 112(10):1734–1746, 2005.

[141] J. Wynbrandt and M. Ludman. *The Encyclopedia of Genetic Disorders and Birth Defects*. Facts on File Library of Health and Living. Facts On File, 2009.

[142] G. Wyszecki and W. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulae*. John Wiley & Sons, 2000.

[143] C. Xu and J. L. Prince. Snakes, shapes, and gradient vector flow. *IEEE Transaction on Image Processing*, 7(3):359–369, 1998.

[144] Y. Yang, T. Wang, N. C. Biswal, X. Wang, M. Sanders, M. Brewer, and Q. Zhu. Optical scattering coefficient estimated by optical coherence tomography correlates with collagen content in ovarian tissue. *Journal of Biomedical Optics*, 16(9):090504–090504, 2011.

[145] M. Yanoff, J. Duker, and J. Augsburger. *Ophthalmology*. An expert consult title online + print. Mosby Elsevier, 2009.

[146] A. L. Yarbus. *Eye Movements and Vision*. Plenum. New York., 1967.

[147] Y. Zhang, M. Brady, and S. Smith. Segmentation of brain mr images through a hidden markov random field model and the expectation-maximization algorithm. *IEEE Transactions on Medical Imaging*, 20(1):45–57, 2001.

[148] A. Zhivov, O. Stachs, J. Stave, and R. F. Guthoff. In vivo three-dimensional confocal laser scanning microscopy of corneal surface and epithelium. *British Journal of Ophthalmology*, 93(5):667–672, 2009.

[149] B. Zuber, L. Stark, and G. Cook. Microsaccades and the velocity-amplitude relationship for saccadic eye movements. *Science*, 150(3702):1459–1460, 1965.

[150] K. Zuiderveld. Contrast limited adaptive histogram equalization. In *Graphics Gems IV*, pages 474–485. Academic Press Professional, Inc., 1994.

Appendix A

Pseudocode

This chapter provides pseudocode as reference implementations for algorithms found throughout the document. Many of these functions use general function names, such as *GaussianBlur*, *Mean*, or *Sqrt*, to indicate an arbitrary implementation. In other cases, these functions use Matlab functions, such as *linspace*, *polyfit*, or *polyval*, using Matlab syntax to communicate the function purpose.

Algorithm 1 (page 170)
> implements boundary localization for a 2D tomogram, T_i, allowing Υ to be determined from a single tomogram.

Algorithm 2 (page 171)
> The structural parameters for a spherical model of a corneal layer, the radius, ρ_ι, and center, c_ι, can be estimated from a set of points, Υ_ι, that exist on the corneal layer using an implementation from Lukacs [75].

Algorithm 3 (page 171)
> The scattering potential random field, Z_2, is estimated using weight least-squares.

Algorithm 4 (page 172)
> Corneal reconstruction without explicitly accounting for motion is demonstrated using Algorithm 1, Algorithm 2, and Algorithm 3.

Algorithm 5 (page 173)
> Pose estimation for each tomogram allows for motion during imaging to be corrected during the reconstruction. Each tomogram is rotated and translated to minimize

the overall reconstruction error, accounting for both corneal structure and scattering potential. The simplified implementation presented here does not include homotopy as it greatly complicates the code structure. For simplicity the calculation of residual error has been reduced to function calls, *ErrorProjectionStructureToTomogram* and *ErrorProjectionScatteringToTomogram*.

Algorithm 1 Corneal Layer Boundary Localization
---
function $\text{LOCATEBOUNDARIES}(T_i)$

    $B \leftarrow \text{GAUSSIANBLUR}(T_i)$ $\triangleright$ Preprocessing

    $E \leftarrow \text{EDGEDETECTION}(B)$ $\triangleright$ Segment edges

    $y_0 \leftarrow \text{GETUPPEREDGEPOINTS}(E)$ $\triangleright$ Assign edges to outer layers

    $y_1 \leftarrow \text{GETLOWEREDGEPOINTS}(E)$

    $s_0 \leftarrow \text{LINSPACE}(-200, 1200, 500)$ $\triangleright$ Space is scaled between -200 and 1200 units in lateral tomogram direction

    $s_1 \leftarrow \text{LINSPACE}(-200, 1200, 500)$

    $\mathcal{C}(s, \hat{\Theta}_0) \leftarrow \text{POLYFIT}(s_0, y_0, 4)$ $\triangleright$ Fit points to curves

    $\mathcal{C}(s, \hat{\Theta}_1) \leftarrow \text{POLYFIT}(s_1, y_1, 4)$

    $s_{0,1} \leftarrow []$ $\triangleright$ Find $s_{0,1}$, shortest distance correspondence between

    for each $s \in s_0$ **do** $\triangleright$ each point on $\mathcal{C}(s, \hat{\Theta}_0)$ to each point on $\mathcal{C}(s, \hat{\Theta}_1)$.

        $s' \leftarrow \arg_{s'} \min \mathcal{C}(s = s, \hat{\Theta}_0) - \mathcal{C}(s = s', \hat{\Theta}_1)$

        $s_{0,1} \leftarrow \text{APPEND}(s_{0,1}, s')$

    end for

    $s_{\text{map}} \leftarrow \text{POLYFIT}(s_0, s_{0,1}, 3)$ $\triangleright$ Arbitrary correspondence function

$$p_{\text{map}}(s_0, \alpha) \leftarrow [x, y]_{map} \leftarrow \alpha \begin{bmatrix} \text{map}(s_0) \\ \mathcal{C}(s = \text{map}(s_0), \hat{\Theta}_1) \end{bmatrix} + (1 - \alpha) \begin{bmatrix} s_0 \\ \mathcal{C}(s = s_0, \hat{\Theta}_0) \end{bmatrix}$$

    $f(\alpha) = []$ $\triangleright$ Pixel intensity integral in α direction

    for each $\alpha \in \text{LINSPACE}(-0.1, 1.1, 500)$ **do**

        $[x, y] = p_{\text{map}}(s_0, \alpha)$

        $f(\alpha) = T_i(x, y)$

    end for

    $W_n \leftarrow \text{ONES}(5, 1)$ $\triangleright$ Use neighbourhoods to smooth out $f(\alpha)$.

    $\mathcal{N}_{n,\alpha} \leftarrow \text{FILTER}(W_n, 1, f(\alpha))$

    $W_f \leftarrow [\text{ONES}(5, 1); \text{ZEROES}(5, 1); \text{ONES}(5, 1)]$ $\triangleright$ Remove near-neighbourhood

    $\mathcal{N}_{f,\alpha} \leftarrow \text{FILTER}(W_f, 1, f(\alpha))$

    for each $i \in 1 : 5$ **do** $\triangleright$ Find local minima

        $h \leftarrow \text{PRIORSPATIALFILTER}(i)$ $\triangleright$ Gaussian prior for α with a mean and std. dev. for each layer boundary from statistics

        $\hat{\alpha}_\iota \leftarrow \text{ARGMIN}(h(\mathcal{N}_{n,\alpha_\iota} - \mathcal{N}_{f,\alpha_\iota}))$

        $[x, y] = p_{\text{map}}(s_0, \alpha_\iota)$

        $\Upsilon_\iota \leftarrow [x, y]$

    end for

    return $[\Upsilon, p_{\text{map}}(s_0, \alpha)]$

end function

Algorithm 2 Spherical Parameter Estimation of Z_1

---

function ESTIMATESPHERE(Υ_ι)

    $\mu_x \leftarrow$ MEAN($\Upsilon_{\iota,x}$)

    $\mu_y \leftarrow$ MEAN($\Upsilon_{\iota,y}$)

    $\mu_z \leftarrow$ MEAN($\Upsilon_{\iota,z}$)

    $d2 \leftarrow \Upsilon_{\iota,x} \hat{.}\, 2 + \Upsilon_{\iota,y} \hat{.}\, 2 + \Upsilon_{\iota,z} \hat{.}\, 2$

    $A \leftarrow \begin{bmatrix} \Upsilon_{\iota,x} .* \Upsilon_{\iota,x} - \mu_x & 2 * \Upsilon_{\iota,x} .* \Upsilon_{\iota,y} - \mu_y & 2 * \Upsilon_{\iota,x} .* \Upsilon_{\iota,z} - \mu_z \\ 0 & \Upsilon_{\iota,y} .* \Upsilon_{\iota,y} - \mu_y & 2 * \Upsilon_{\iota,y} .* \Upsilon_{\iota,z} - \mu_z \\ 0 & 0 & \Upsilon_{\iota,z} .* \Upsilon_{\iota,z} - \mu_z \end{bmatrix}$

    $A \leftarrow A + A^T$

    $B \leftarrow$ ZEROS($3, 1$)

    $B_1 \leftarrow$ MEAN($d2 .* \left(\Upsilon_{\iota,x} - \mu_x\right)$)

    $B_2 \leftarrow$ MEAN($d2 .* \left(\Upsilon_{\iota,y} - \mu_y\right)$)

    $B_3 \leftarrow$ MEAN($d2 .* \left(\Upsilon_{\iota,z} - \mu_z\right)$)

    $c_\iota \leftarrow$ SOLVE($Ac_\iota = B$)

    $d2 \leftarrow \left(\left(\Upsilon_{\iota,x} - \mu_x\right) \hat{.}\, 2 + \left(\Upsilon_{\iota,y} - \mu_y\right) \hat{.}\, 2 + \left(\Upsilon_{\iota,z} - \mu_z\right) \hat{.}\, 2\right)$

    $\mu_{d2} \leftarrow$ MEAN($d2$)

    $\rho_\iota \leftarrow$ SQRT(μ_{d2})

    return $[\rho_\iota, c_\iota]$

end function

Algorithm 3 Weighted Least-Squares Parameter Estimation of Z_2

---

function ESTIMATEZ2($\mathbf{m}, R, C_2, \wp_2$)

    $\mathbf{m}_2 \leftarrow \wp_2\big(\mathbf{m}\big)$

    $\mu_2 \leftarrow$ MEAN($\mathbf{m}_2$)

    $U \leftarrow$ CHOLESKYDECOMPOSITION($C_2^T R^{-1} C_2 + Q$) $\triangleright\ U^T U = C_2^T R^{-1} C_2 + Q$

    $X_1 \leftarrow$ SOLVE($U^T X_1 = C_2^T R^{-1}\big(\mathbf{m}_2 - \mu_2\big)$) $\triangleright$ Solve using forward substitution

    $\hat{Z}_2 \leftarrow$ SOLVE($U\hat{Z}_2 = X_1$) $\triangleright$ Solve using backward substitution

    return $\hat{Z}_2$

end function

Algorithm 4 Reconstruction of Z without Motion Correction

---

function RECONSTRUCT(T, R, Θ)
    $\mathbf{m} \leftarrow T_.$
    $C_1 \leftarrow []$ ▷ Determine structural correspondence
    for all $i \in T_i$ **do**
        $C_i \leftarrow$ GETCORRESPONDENCE(Θ_i)
        $C_1 \leftarrow$ APPEND(C_1, C_i)
    end for
    $\Upsilon \leftarrow []$ ▷ Locate corneal boundaries in each tomogram
    for all $i \in T_i$ **do**
        $[U, p_{i,\text{map}}(s_0, \alpha)] \leftarrow$ LOCATEBOUNDARIES(T_i)
        $\Upsilon \leftarrow$ APPEND(Υ, U)
    end for
    for each $\iota \in 1 : 5$ **do**   ▷ Estimate structure for each boundary
        $[\rho_\iota, c_\iota] \leftarrow$ ESTIMATESPHERE(Υ_ι)
    end for
    $C_2 \leftarrow []$ ▷ map space from tomogram to $[s, \alpha]$
    $\wp_2(\mathbf{m}) \leftarrow \mathbf{m}$ ▷ For this implementation, all preprocessing is moved into C_2
    $[s_0\alpha] \leftarrow$ NDGRID$(-200 : 1200, -.1 : 1.1)$
    for all $i \in T_i$ **do**
        $[xy] \leftarrow p_{i,\text{map}}(s_0, \alpha)$
        $s_{\text{map}} \leftarrow$ TRISCATTEREDINTERP(x, y, s_0)
        $\alpha_{\text{map}} \leftarrow$ TRISCATTEREDINTERP(x, y, α)
        $[xyz] \leftarrow$ GETCOORDINATES$(\mathbf{m}_i)$
        $[sa] \leftarrow [s_{\text{map}}(x, y), \alpha_{\text{map}}(x, y)]$
        $C_{2,i}$ GETCORRESPONDENCE$(x, y, z, s, a, \Theta_i)$
        $C_2 \leftarrow$ APPEND$(C_2, C_{2,i})$
    end for
    $\hat{Z}_2 \leftarrow$ ESTIMATEZ2$(\mathbf{m}, R, C_2, \wp_2)$
    return $[\rho, c, \hat{Z}_2]$
end function

Algorithm 5 Motion Corrected Parameter Estimation of Z

---
function ESTIMATEPOSE(T, R, Θ_0) ▷ Create initial reconstruction based on Θ_0
     $\Theta \leftarrow \Theta_0$
     for each $l \in 1:5$ **do**   ▷ Run for five iterations
         $[\rho, c, \hat{Z}_2] \leftarrow$RECONSTRUCT$(T, R, \Theta)$
         $f(\Theta) \leftarrow$@ESTIMATEPOSEOPTFUN(T, R, Θ)
         $e, \hat{\Theta} \leftarrow$ARGMINTHETA$(f(\Theta))$ ▷ Solve using optimization implementation
     end for
     return $\hat{\Theta}$
end function
function ESTIMATEPOSEOPTFUN$(T, R, \Theta, \rho, c, \hat{Z}_2)$
     $e \leftarrow 0$
     $w \leftarrow .5$ ▷ Weight between structure and scattering potential
     for all $i \in T_i$ **do**
         $e_{i,\text{structure}} \leftarrow$ERRORPROJECTIONSTRUCTURETOTOMOGRAM(T_i, ρ, c)
         $e_{i,\text{scattering}} \leftarrow$ERRORPROJECTIONSCATTERINGTOTOMOGRAM$(T_i, \hat{Z}_2)$
         $e \leftarrow e + we_{i,\text{structure}} + (1 - w)e_{i,\text{scattering}}$
     end for
     return e ▷ return residual error
end function